A Survey of the Skolem and Positivity Problems for Linear Recurrence Sequences

Piotr Bacik[1,2], Toghrul Karimov[2], Florian Luca[3,2], Joris Nieuwveld[1], Joël Ouaknine[2], David Purser[4], and James Worrell[1]

[1]Department of Computer Science, Oxford University, UK
[2]Max Planck Institute for Software Systems, Saarland Informatics Campus, Germany
[3]Mathematics Division, Stellenbosch University, South Africa
[4] School of Computer Science and Informatics, University of Liverpool, UK

July 26, 2026

Abstract

The Skolem Problem asks to determine whether a given linear recurrence sequence (LRS) over the integers contains a zero, while the Positivity Problem asks whether all of its terms are non-negative. The decidability of both problems has been open for many decades, despite the ubiquity of LRS and the abundance of powerful mathematical tools for their analysis. We survey the state of the art on these two problems and their natural variants, including the Bi-Skolem and Ultimate Positivity Problems. We describe known decidability results for low-order sequences, as well as conditional decidability results predicated on classical number-theoretic conjectures (most notably the Exponential Local-Global Principle, the p-adic Schanuel Conjecture, and Cramér-type heuristics on prime gaps). We pay particular attention to recent developments, including complexity advances placing the Skolem Problem at low orders in randomized polynomial time, the construction of Universal Skolem Sets of density one, algorithmic developments around the SKOLEM tool, and Diophantine hardness barriers for the Positivity Problem beyond order 5. We also exhibit two concrete integer LRS, of orders 6 and 10, for which we are currently unable to solve the Skolem and Positivity Problems respectively. We close with a *tour d'horizon* of related problems and applications, including the Skolem Problem in positive characteristic, continuous-time analogues, robust and pseudo-variants, the role of Skolem and Positivity as benchmarks for hardness in the algorithmic verification of loop termination and probabilistic systems, and a list of what we view as the central open problems in the area.

1 Introduction

Linear recurrence sequences, of which the Fibonacci numbers are the most emblematic example, arise in a diverse range of contexts, particularly in mathematics and computer science, but also in fields as varied as physics, biology, and economics. From a computational standpoint, a number of decision problems concerning linear recurrence sequences have come to prominence over the last four decades, with applications to automata theory, formal languages and power series, loop termination, stochastic verification, and control theory, among many other areas. Notwithstanding the apparent simplicity of linear recurrence sequences, the decidability of several of the most natural and compelling algorithmic problems on this class of objects remains stubbornly open. It is striking that even the relatively limited progress that has been made in the field has relied on some of the most celebrated theorems of modern number theory.

The present survey treats two such problems—the *Skolem Problem* and the *Positivity Problem*—together with certain natural variants and a number of closely related questions. We trace

the historical development of the field, present the state of the art, and survey advances made over the last decade or so. We assume on the part of the reader some familiarity with the standard background of theoretical computer science and discrete mathematics; the more specialized number-theoretic and p-adic prerequisites are developed in Section 2.

Let us begin with some definitions. A *rational linear recurrence sequence* (LRS) is an infinite sequence $\boldsymbol{u} = \langle u_n \rangle_{n=0}^{\infty}$ of rational numbers satisfying a recurrence relation of the form

$$u_{n+d} \;=\; a_{d-1}\, u_{n+d-1} \;+\; \cdots \;+\; a_0\, u_n \qquad (n \in \mathbb{N}), \tag{1}$$

together with prescribed initial values $u_0, u_1, \ldots, u_{d-1} \in \mathbb{Q}$. We assume throughout that $a_0 \neq 0$, which entails that the recurrence can be run not only forward but also backward from the initial values, yielding a bi-infinite sequence to which we shall return shortly. The smallest d for which $\boldsymbol{u}$ obeys a relation of the form above is called the *order* of the sequence. We associate with the recurrence relation (1) its *characteristic polynomial*

$$p(x) \;=\; x^d - a_{d-1}\, x^{d-1} - \cdots - a_1\, x - a_0;$$

the characteristic polynomial of an LRS is then defined to be that of its minimal-order recurrence relation. An LRS is said to be *simple* if its characteristic polynomial has exclusively simple (i.e., non-repeated) roots.

Every LRS admits a closed-form representation as an *exponential polynomial*: if $\lambda_1, \ldots, \lambda_m$ are the distinct characteristic roots of $\boldsymbol{u}$, then there exist univariate polynomials $Q_1, \ldots, Q_m$ with algebraic coefficients such that

$$u_n \;=\; \sum_{j=1}^{m} Q_j(n)\, \lambda_j^n \qquad (n \in \mathbb{N}). \tag{2}$$

One can show that an LRS is simple if and only if all of the Q_j are constant polynomials; we postpone the proof to Section 3.

Example 1. The Fibonacci sequence $\langle 0, 1, 1, 2, 3, 5, 8, 13, \ldots \rangle$ satisfies the recurrence $f_{n+2} = f_{n+1} + f_n$ with $f_0 = 0$ and $f_1 = 1$. Its characteristic polynomial $x^2 - x - 1$ has as its roots the golden ratio $\varphi = (1+\sqrt{5})/2$ together with its Galois conjugate $\tilde{\varphi} = (1-\sqrt{5})/2$, and the Fibonacci sequence is therefore simple. It admits the well-known exponential-polynomial representation

$$f_n \;=\; \frac{1}{\sqrt{5}}\, \varphi^n \;-\; \frac{1}{\sqrt{5}}\, \tilde{\varphi}^{\,n},$$

known as *Binet's formula*.

The first protagonist of the present survey can now be introduced.

Problem 2 (The Skolem Problem)**.** *Given an LRS $\boldsymbol{u}$, does there exist $n \in \mathbb{N}$ such that $u_n = 0$?*

The decidability of the Skolem Problem has been open for close to a century [58], a state of affairs described as "faintly outrageous" by T. Tao [120]—"it is saying that we do not know how to decide the halting problem even for 'linear' automata"—and as a "mathematical embarrassment" by R. J. Lipton and K. W. Regan [77].

The second protagonist is the following.

Problem 3 (The Positivity Problem)**.** *Given an LRS $\boldsymbol{u}$, is it the case that $u_n \geq 0$ for all $n \in \mathbb{N}$?*

The decidability of the Positivity Problem is stated as open in literature going back at least to the 1970s [117, 75, 25], as well as in many more recent treatments [57, 22, 72, 121, 99]. The Positivity Problem is at least as hard as the Skolem Problem. The folklore reduction relies on closure properties of the class of LRS (see Proposition 18). First, the Skolem Problem on

rational sequences can be reduced to the version on integer sequences by scaling: given an LRS $\boldsymbol{u} := \langle u_n \rangle_{n=0}^{\infty}$, let L be the least common denominator of the rational numbers appearing in the recurrence and initial conditions; the sequence $\langle L^{n+1} u_n \rangle_{n=0}^{\infty}$ is then integer-valued and has the same zeros as $\boldsymbol{u}$. Next, given an integer LRS $\boldsymbol{v} := \langle v_n \rangle_{n=0}^{\infty}$, the Hadamard square $\langle v_n^2 \rangle_{n=0}^{\infty}$ is itself an LRS (of order at most $\binom{d+1}{2}$, where d is the order of $\boldsymbol{v}$), and $\boldsymbol{v}$ has no zero if and only if $\langle v_n^2 - 1 \rangle_{n=0}^{\infty}$ is non-negative. Decidability of Positivity would therefore entail decidability of Skolem. There is reason to believe, moreover, that the Positivity Problem is in some sense substantially harder than the Skolem Problem; we develop this point in Section 7.

A natural variant of the Positivity Problem is the following.

Problem 4 (The Ultimate Positivity Problem)**.** *Given an LRS $\boldsymbol{u}$, does there exist $N \in \mathbb{N}$ such that $u_n \geq 0$ for all $n \geq N$?*

While the Skolem and Positivity Problems remain open for the class of simple LRS, as we shall see below, the Ultimate Positivity Problem is decidable on this class. It is worth remarking that, in a natural measure-theoretic sense, 'most' LRS are simple: for instance, when drawing a polynomial uniformly at random by sampling its coefficients independently and uniformly at random from a large interval, the probability that the polynomial has simple roots tends to one as the length of the interval grows.

The Skolem–Mahler–Lech Theorem. The structure of the zero set $\{n \in \mathbb{N} : u_n = 0\}$ of an LRS is described by the celebrated theorem of T. Skolem, K. Mahler, and C. Lech [115, 84, 73].

Theorem 5 (Skolem–Mahler–Lech)**.** *Given a linear recurrence sequence $\boldsymbol{u}$, the set $\{n \in \mathbb{N} : u_n = 0\}$ is a union of finitely many full arithmetic progressions, together with a finite set.*

This theorem will be discussed at length in Sections 3 and 5; for now we make just three remarks. Firstly, although this survey focuses on LRS over the rational numbers, the Skolem–Mahler–Lech Theorem applies more generally to LRS over fields of characteristic zero. Secondly, the statement of the theorem can be refined by considering the notion of *non-degeneracy* of an LRS. An LRS is non-degenerate provided that no quotient of two distinct characteristic roots is a root of unity. A given LRS can be effectively decomposed as an interleaving of finitely many non-degenerate subsequences, some of which may be identically zero. The core of the Skolem–Mahler–Lech Theorem is then the assertion that a non-zero non-degenerate LRS has only finitely many zeros. Unfortunately, all known proofs of this last assertion are ineffective: it is not known how to compute the finite set of zeros of a given non-degenerate linear recurrence sequence. Existence of a procedure to do so is readily seen to be equivalent to decidability of the Skolem Problem. Finally, there is no similar simple structural characterization of the positivity set $\{n \in \mathbb{N} : u_n \geq 0\}$ of an LRS. What is known in general is that the positivity set always has a well-defined density [22], which can in fact be computed to arbitrary precision [67].

For the purposes of this survey, the effective nature of the decomposition of an arbitrary LRS as an interleaving of finitely many non-degenerate sequences allows us to assume without loss of generality, when convenient, that all our LRS are non-degenerate, and hence are either identically zero or possess only finitely many zeros.

Bi-infinite sequences and the Bi-Skolem Problem. A useful concept in the study of LRS is that of their bi-infinite completion [31]. Since we have assumed that the tail coefficient a_0 of the recurrence (1) is non-zero, the recurrence can be run backward as well as forward from the initial values, yielding a *linear recurrence bi-sequence* (LRBS) $\boldsymbol{u} = \langle u_n \rangle_{n=-\infty}^{\infty}$. The bi-infinite completion of the Fibonacci sequence, for example, is

$$\langle \ldots, 5, -3, 2, -1, 1, 0, 1, 1, 2, 3, 5, \ldots \rangle.$$

The notions of order, simplicity, and non-degeneracy all carry over in the obvious way to LRBS, and the exponential-polynomial form (2) extends without change to all integers $n \in \mathbb{Z}$. The Skolem–Mahler–Lech Theorem also remains valid in the bi-infinite setting: a non-degenerate LRBS has only finitely many zeros. The natural analogue of the Skolem Problem in this setting is the following.

Problem 6 (The Bi-Skolem Problem)**.** *Given an LRBS* $\boldsymbol{u}$*, does there exist* $n \in \mathbb{Z}$ *such that* $u_n = 0$*?*

A central motivation for considering the Bi-Skolem Problem is the existence of the *Exponential Local-Global Principle*, a classical conjecture that we discuss in detail in Section 6. Informally, restricted to our setting, the principle asserts that a simple LRBS fails to have a zero if and only if this absence of zeros is witnessed by a suitable congruence: that is, there exists an integer $m \geq 2$ such that, viewed in $\mathbb{Z}/m\mathbb{Z}$, the LRBS is well defined and everywhere non-zero. The truth of this conjectured principle would immediately yield an algorithm for the Bi-Skolem Problem on simple LRBS, by parallel search for either a zero of the sequence or a witnessing congruence. Verifying a candidate congruence is effective since the resulting bi-sequence of residue classes is periodic, thanks to the pigeonhole principle.

It is worth noting in passing that the Exponential Local-Global Principle essentially concerns bi-sequences and fails for ordinary LRS. Consider, for example, the shifted Fibonacci sequence $\langle 1, 1, 2, 3, 5, 8, 13, \ldots \rangle$. This LRS is zero-free, but its bi-infinite completion is the same as that of the true Fibonacci sequence and therefore contains a zero. Moreover the sequence is periodic when interpreted in $\mathbb{Z}/m\mathbb{Z}$ for any non-zero modulus m. Thus the shifted Fibonacci sequence necessarily contains infinitely many zeros modulo m, for any $m \geq 2$. Metaphorically speaking, the shifted Fibonacci sequence is haunted by the 'ghost' of a zero in its past.

Where Skolem and Positivity arise. The Skolem and Positivity Problems sit at a striking junction of several mathematical disciplines: algebraic number theory, Diophantine approximation, p-adic analysis, automata and formal language theory, and algorithmic verification. Beyond their intrinsic interest, they serve as central reference points for a wide range of decision problems in theoretical computer science, where they appear either as the algorithmic question of interest or as a benchmark of hardness. A non-exhaustive sampler includes the following:

- *Theoretical biology and developmental systems* [75, 117], where Skolem-type questions arise in the analysis of Lindenmayer systems and their growth functions.
- *Formal power series and weighted automata* [26, 107, 108], where the Skolem Problem arises in connection with decision problems on the support of rational power series.
- *Loop termination and program verification* [97, 101, 65], where reachability of zero (or of a negative number) by an LRS captures the termination of simple linear loops over the integers.
- *Reachability for linear dynamical systems* [61, 41, 15], encompassing the Orbit Problem, the Polyhedron-Hitting Problem, and their parametric and control-theoretic analogues.
- *Markov chains, Markov decision processes, and probabilistic programs* [2, 122, 1, 103, 19], where Skolem-hardness and Positivity-hardness arise pervasively in quantitative verification.
- *Control theory* [35, 54], where the Skolem Problem appears in problems of stability, observability, and controllability of linear systems.
- *Matrix semigroups* [24, 38], where the Skolem Problem corresponds to a one-dimensional mortality-type question for a single matrix.

In many of these areas, natural decision problems are equivalent to, or at least as hard as, the Skolem and Positivity Problems.

The contemporary landscape. Progress on the Skolem and Positivity Problems has been limited, with new techniques typically yielding decidability for sequences subject to fairly stringent restrictions on order or on the structure of the characteristic roots. The unconditional state of the art for the Skolem Problem stood essentially still from the mid-1980s until very recently, with decidability known only for LRS of order at most 4, established through the work of M. Mignotte, T. N. Shorey, and R. Tijdeman [91] and of N. K. Vereshchagin [125], both leveraging Baker's celebrated theorem on linear forms in logarithms [17]. Decidability of Skolem at order 5 remains open. The state of the art for the Positivity Problem is similarly limited to low-order sequences: decidability is known only for LRS of order at most 5 [99] (and, for simple LRS, of order at most 9 [98]), and a deeper analysis reveals that establishing decidability of Positivity at order 6 would entail substantial new results in Diophantine approximation, specifically concerning how well certain transcendental numbers can be approximated by rationals of bounded height.

The last decade has nonetheless witnessed a flurry of activity, yielding both unconditional advances and a rich array of conditional results predicated on classical number-theoretic conjectures. We have seen, on the unconditional side, completion of the picture for the Skolem Problem at order 4 over the algebraic numbers [11]; a substantial sharpening of the complexity of the Skolem Problem at low orders, placing the problem in coRP for LRS of order at most 4 [13], improving the previously best known $\mathsf{NP}^{\mathsf{RP}}$ bound by a full exponential; the construction of *Universal Skolem Sets* [79, 81, 78] of progressively higher density; and an unconditional Universal Skolem Set of density one [83]. On the conditional side, we can mention the decidability of the Skolem Problem for simple LRS assuming the Exponential Local-Global Principle and a weak form of the p-adic Schanuel Conjecture [31];[1] new results on p-adic variants of the Skolem Problem, as well as decidability of the Simultaneous Skolem Problem (determining whether two coprime LRS have a common zero), in both cases assuming the p-adic Schanuel Conjecture [12]; and decidability of the Skolem Problem over all LRS, assuming Cramér-type heuristics on prime gaps [82, 83].

Scope and structure of this survey. Our aim is to survey the present state of the Skolem and Positivity Problems, with an emphasis on results and techniques that have appeared since the mid-2010s.[2] We focus throughout on rational and algebraic LRS, i.e., on rings and fields of characteristic 0. The complementary setting of positive characteristic is much better understood than the characteristic-0 case, owing in particular to the work of H. Derksen [49], and we shall not treat it here; we return to the positive-characteristic story in Section 9 as part of a broader tour of related problems.

The structure of the survey is as follows. Section 2 introduces the mathematical background needed throughout: algebraic numbers and heights, Baker's Theorem on linear forms in logarithms, and the fundamentals of p-adic analysis. Section 3 then develops the basic theory of linear recurrence sequences in characteristic 0, including their exponential-polynomial form, non-degeneracy, and closure properties.

Section 4 introduces what we view as the unifying conceptual framework of the survey, namely the notion of a *certificate* of zero-freeness for the Skolem Problem (and of non-negativity for the Positivity Problem). All known unconditional decidability results, and most known conditional

[1] The weak p-adic Schanuel Conjecture is not required for simple LRS of order 5 [76], nor in fact for simple LRS of orders 6 and 7 [95].

[2] An earlier and considerably shorter exposition of some of the material below, tracing the development of the field up to that point, appeared in 2015 as a SIGLOG News article [101]; the present survey substantially supersedes and extends that exposition.

ones, can be cast in terms of three classes of such certificates: separation bounds (in either Archimedean or non-Archimedean form), modular invariants, and semialgebraic invariants. The remainder of the survey is, to a significant extent, organized around which class of certificate is being deployed.

Sections 5 and 6 are the first core technical block, devoted to the Skolem Problem. Section 5 treats the unconditional state of the art—the Skolem–Mahler–Lech Theorem, the unconditional decidability results at low orders and the MSTV class, and recent complexity advances [13]. Section 6 surveys conditional and alternative approaches: the Bi-Skolem Problem and the Exponential Local-Global Principle, the p-adic Schanuel Conjecture and the conditional decidability results for simple LRS, the p-adic Skolem Problem, the construction of Universal Skolem Sets, the Cramér heuristic for large zeros, and twisted rational zeros.

Section 7 forms the second core technical block, treating the Positivity and Ultimate Positivity Problems: the unconditional decidability results at low orders and on simple LRS, the Diophantine hardness barrier at order 6 for general LRS, the frontier at order 10 for simple LRS, and some of the related reduction machinery around the Positivity Problem.

Section 8 is devoted to a small collection of worked examples, intended to make the abstract obstacles to decidability more concrete by exhibiting specific LRS at the boundary of what is currently known, and to illustrate the certificate framework in action. In particular, Section 8 contains two concrete instances of integer LRS for which we do not know how to solve the Skolem and Positivity Problems respectively. Section 9 surveys the broader context: the Skolem Problem in positive characteristic, continuous-time analogues, robust and pseudo-variants, and a list of open problems.

What we do not cover. A full treatment of the connections between the Skolem and Positivity Problems and algorithmic verification of linear and probabilistic systems would require a survey of its own. We therefore content ourselves in Section 9 with pointers to the literature. Likewise, we shall not discuss continuous-time linear dynamical systems beyond a brief mention. Finally, we omit any treatment of holonomic (P-finite) sequences, which generalize linear recurrences with constant coefficients to the setting of polynomial coefficients; for an entry point into that literature we refer to [66, 96].

2 Mathematical Preliminaries

In this section we collect the number-theoretic and p-adic background needed in the remainder of the survey. Our treatment is deliberately concise: our purpose is to fix notation and to introduce sufficient machinery to enable the developments that follow. For thorough expositions of algebraic number theory, the reader is referred to [71, 94]; for transcendence theory and Baker's Theorem, to [17, 18]; for p-adic analysis, to [69, 40, 106]. A reader already comfortable with these topics may safely skim this section on a first reading and refer back as needed.

2.1 Algebraic Numbers, Number Fields, and Heights

A complex number $\alpha \in \mathbb{C}$ is *algebraic* if it is a root of some non-zero polynomial with rational coefficients. Among all such polynomials, there is a unique monic polynomial of minimal degree, called the *minimal polynomial* of α over $\mathbb{Q}$ and denoted $\mu_\alpha(x) \in \mathbb{Q}[x]$. The degree of μ_α is called the *degree* of α. The other roots of μ_α are called the *(Galois) conjugates* of α. An algebraic number is an *algebraic integer* if its minimal polynomial has integer coefficients. The set of all algebraic numbers forms a field, denoted $\overline{\mathbb{Q}}$, and the algebraic integers form a subring thereof.

Let K be a field extension of $\mathbb{Q}$. The dimension $[K : \mathbb{Q}]$ of K as a $\mathbb{Q}$-vector space is called its *degree*. A finite-degree field extension is called a *number field*. Every number field is of the form $K = \mathbb{Q}(\alpha)$ for some algebraic number α, in which case $[K : \mathbb{Q}]$ equals the degree of α. The

set of algebraic integers contained in K forms a subring $\mathcal{O}_K$ of K, called the *ring of integers* of K.

For computational purposes, an algebraic number α of degree d is conventionally represented by its minimal polynomial $\mu_\alpha(x) \in \mathbb{Q}[x]$ together with sufficiently accurate rational approximations of its real and imaginary parts to distinguish it from its conjugates (an accuracy polynomial in the degree and height of α, in the sense defined below, suffices). With such a representation, all standard arithmetic operations on algebraic numbers, and decision of equality, can be carried out effectively in time polynomial in the bit-size of the input; see [46] for a comprehensive account.

Absolute values. Let K be a number field. An *absolute value* on K is a function $|\cdot| : K \to \mathbb{R}_{\geq 0}$ satisfying $|x| = 0$ if and only if $x = 0$, $|xy| = |x| \cdot |y|$, and $|x+y| \leq |x| + |y|$ for all $x, y \in K$. Two absolute values are *equivalent* if they induce the same topology on K. An equivalence class of absolute values on K is called a *place* of K. The places of K are of two kinds.

The *Archimedean places* (or *infinite places*) of K arise from the $[K : \mathbb{Q}]$ distinct field embeddings of K in $\mathbb{C}$. Each such place gives rise to an absolute value $|x|_\sigma := |\sigma(x)|$, where on the right the standard absolute value on $\mathbb{C}$ is used. Hence complex-conjugate embeddings give rise to the same place.

The *non-Archimedean places* (or *finite places*) of K correspond to the non-zero prime ideals $\mathfrak{p}$ of $\mathcal{O}_K$. Each such place arises as follows. Every non-zero prime ideal $\mathfrak{p}$ of $\mathcal{O}_K$ lies over a unique rational prime p, in the sense that $\mathfrak{p} \cap \mathbb{Z} = p\mathbb{Z}$. There is associated with $\mathfrak{p}$ a discrete valuation $v_\mathfrak{p} : K^\times \to \mathbb{Z}$, defined for $x \in \mathcal{O}_K \setminus \{0\}$ to be the largest integer k such that $x \in \mathfrak{p}^k$, and extended to $K^\times$ multiplicatively; one further sets $v_\mathfrak{p}(0) := +\infty$. The corresponding absolute value is given by $|x|_\mathfrak{p} := p^{-v_\mathfrak{p}(x) f_\mathfrak{p} / e_\mathfrak{p}}$, where $e_\mathfrak{p}$ and $f_\mathfrak{p}$ are respectively the *ramification index* and *residue degree* of $\mathfrak{p}$ over p; for our purposes the precise normalization will rarely matter, and we shall be content to write $v_p(\cdot)$ when working over $\mathbb{Q}$ itself.

A defining feature of non-Archimedean absolute values is the so-called *ultrametric inequality*, which strengthens the usual triangle inequality:

$$|x + y|_\mathfrak{p} \;\leq\; \max\bigl(|x|_\mathfrak{p},\, |y|_\mathfrak{p}\bigr).$$

Heights. The notion of *height* provides a quantitative measure of the arithmetic complexity of an algebraic number. For a rational number a/b in lowest terms, with $b > 0$, the natural notion of size is $\max(|a|, b)$. Heights generalize this to arbitrary algebraic numbers in a way that interacts well with field operations.

The *absolute (logarithmic) Weil height* of an algebraic number α of degree d, with conjugates $\alpha = \alpha_1, \alpha_2, \ldots, \alpha_d$, is defined as follows. Writing

$$a_d \prod_{i=1}^{d} (x - \alpha_i) \;\in\; \mathbb{Z}[x] \quad (\text{gcd of coefficients } 1,\; a_d > 0)$$

for the primitive integer multiple of the minimal polynomial μ_α, one sets

$$h(\alpha) \;:=\; \frac{1}{d} \left(\log a_d \;+\; \sum_{i=1}^{d} \log \max(1, |\alpha_i|) \right).$$

For a rational number a/b in lowest terms with $b > 0$, this specializes to $h(a/b) = \log \max(|a|, b)$, recovering the naive notion. The Weil height enjoys a number of properties, of which we record only those we shall need:

- $h(\alpha) \geq 0$ for all $\alpha \in \overline{\mathbb{Q}}$, with equality if and only if $\alpha = 0$ or α is a root of unity (a result known as Kronecker's theorem),

- $h(\alpha+\beta) \le h(\alpha)+h(\beta)+\log 2$ and $h(\alpha\beta) \le h(\alpha)+h(\beta)$,
- $h(\alpha^n) = |n| \cdot h(\alpha)$ for all $n \in \mathbb{Z}$,
- $h(\alpha) = h(\beta)$ whenever α and β are Galois conjugates.

For our algorithmic purposes, the key fact is that $h(\alpha)$ is effectively computable from a description of α by its minimal polynomial, and that the bit-size of α in the standard representation and the height $h(\alpha)$ are polynomially related. Precise bounds may be found in [126].

2.2 Baker's Theorem on Linear Forms in Logarithms

A pivotal tool in the algorithmic study of linear recurrence sequences is Baker's celebrated theorem on linear forms in logarithms of algebraic numbers [17], for which A. Baker was awarded the Fields Medal in 1970. Roughly speaking, Baker's Theorem provides effective lower bounds for non-zero linear combinations of logarithms of algebraic numbers with algebraic coefficients. We state below a version sufficient for our applications; for the sharpest known formulations, the reader is referred to [18, 89].

Theorem 7 (Baker). *Let $\alpha_1, \dots, \alpha_n$ be non-zero algebraic numbers, and let $\log\alpha_1, \dots, \log\alpha_n$ denote any choice of complex logarithms. Let $b_1, \dots, b_n \in \mathbb{Z}$, set $B := \max(|b_1|, \dots, |b_n|, 3)$, and let*

$$\Lambda \;:=\; b_1 \log\alpha_1 + b_2 \log\alpha_2 + \cdots + b_n \log\alpha_n.$$

If $\Lambda \neq 0$, then there is an effectively computable constant $C > 0$, depending only on n, on the degree of the number field $\mathbb{Q}(\alpha_1, \dots, \alpha_n)$, and on the heights $h(\alpha_i)$, such that

$$|\Lambda| \;\ge\; B^{-C}.$$

The constant C in Theorem 7 is, in concrete formulations, of the order

$$C \;=\; (16nd)^{2(n+2)} \prod_{i=1}^{n} \max\bigl(h(\alpha_i),\, |\log\alpha_i|/d,\, 1/d\bigr),$$

where d is the degree of the number field $\mathbb{Q}(\alpha_1, \dots, \alpha_n)$. The dependence of this lower bound on n is doubly exponential, but for fixed n the dependence on the remaining parameters is benign. This is precisely the regime of interest for LRS of bounded order, whose associated linear forms involve only a bounded number of logarithms.

The p-adic version. In addition to the classical Archimedean version, a p-adic analogue of Baker's Theorem is needed. We formulate below a statement of this result, established by K. Yu in [127]. The statement is closely parallel to the Archimedean case, albeit phrased in terms of the $\mathfrak{p}$-adic valuation of the quantity $\alpha_1^{b_1} \cdots \alpha_n^{b_n} - 1$ rather than in terms of a linear form in p-adic logarithms.

Theorem 8 (Baker, p-adic version). *Let $\alpha_1, \dots, \alpha_n$ be non-zero algebraic numbers contained in a number field K of degree d. Let $\mathfrak{p} \subseteq \mathcal{O}_K$ be a prime ideal lying above the rational prime p, such that $v_{\mathfrak{p}}(\alpha_i) = 0$ for $i = 1, \dots, n$. Let $b_1, \dots, b_n \in \mathbb{Z}$ be such that*

$$\Xi := \alpha_1^{b_1} \cdots \alpha_n^{b_n} - 1 \neq 0.$$

Then

$$v_{\mathfrak{p}}(\Xi) \;\le\; C\, A_1 \cdots A_n \, \log B,$$

where

$$A_j \ge \max\{\, h(\alpha_j),\, \log p \,\} \quad (j = 1, \dots, n), \qquad B \ge \max\{\, |b_1|, \dots, |b_n|,\, 3 \,\},$$

and $C > 0$ is an effectively computable constant depending only on n, d, and $\mathfrak{p}$.

The Archimedean and p-adic versions of Baker's Theorem are the basic tools underlying virtually all known unconditional decidability results for the Skolem and Positivity Problems at low orders, and we shall invoke them frequently in Sections 5 and 7.

2.3 p-adic Numbers and p-adic Analysis

The p-adic numbers play a central role in the study of the Skolem Problem, beginning with Skolem's original method, which underlies the Skolem–Mahler–Lech Theorem (Section 5.2). We collect here the basic definitions and the technical results we shall need.

Construction. Fix a rational prime p (i.e., a prime number in $\mathbb{N}$). The p-adic absolute value $|\cdot|_p$ on $\mathbb{Q}$ is defined by $|0|_p := 0$ and, for $x = p^k a/b$ with $\gcd(a,p) = \gcd(b,p) = 1$, by $|x|_p := p^{-k}$. The corresponding p-adic valuation is $v_p(x) := -\log_p |x|_p \in \mathbb{Z} \cup \{+\infty\}$. The field $\mathbb{Q}_p$ of *p-adic numbers* is the completion of $\mathbb{Q}$ with respect to $|\cdot|_p$. Its valuation ring

$$\mathbb{Z}_p \;:=\; \{x \in \mathbb{Q}_p : |x|_p \leq 1\}$$

is the ring of *p-adic integers*, and is the natural domain on which the p-adic analytic functions arising from LRS will be defined.

The field $\mathbb{Q}_p$ is locally compact and complete, but not algebraically closed. For our applications, we shall typically need to work in a finite extension $K_\mathfrak{p}$ of $\mathbb{Q}_p$, large enough to contain the characteristic roots of the LRS under consideration. Such an extension comes with its own ring of integers $\mathcal{O}_\mathfrak{p}$, and the p-adic absolute value $|\cdot|_p$ extends uniquely from $\mathbb{Q}_p$ to $K_\mathfrak{p}$. For thoroughness, we mention that the algebraic closure $\overline{\mathbb{Q}_p}$ of $\mathbb{Q}_p$ is itself not complete, and its completion $\mathbb{C}_p$ is both algebraically closed and complete; $\mathbb{C}_p$ is the natural p-adic analogue of the complex numbers. We shall need $\mathbb{C}_p$ only occasionally; it also provides a uniform setting in which the various finite extensions encountered in the sequel may all be embedded.

Power series and discs. For a in (a finite extension of) $\mathbb{Q}_p$ and $r > 0$, write

$$D(a,r) \;:=\; \{x : |x-a|_p \leq r\}$$

for the closed disc of radius r around a. (Ultrametrically, discs are simultaneously open and closed, and every point of a disc is a center: $D(a,r) = D(b,r)$ whenever $|a-b|_p \leq r$.) The ring of integers $\mathbb{Z}_p$ is itself the closed unit disc $D(0,1) \subseteq \mathbb{Q}_p$.

A power series $f(x) = \sum_{n\geq 0} c_n x^n$ with p-adic coefficients c_n converges on $D(0,r)$ provided that $|c_n|_p\, r^n \to 0$ as $n \to \infty$. Standard examples include the p-adic exponential

$$\exp_p(x) \;:=\; \sum_{n\geq 0} \frac{x^n}{n!},$$

which converges on the open disc $\{x : |x|_p < p^{-1/(p-1)}\}$, and the p-adic logarithm

$$\log_p(1+x) \;:=\; \sum_{n\geq 1} \frac{(-1)^{n+1} x^n}{n},$$

which converges on $\{x : |x|_p < 1\}$. The latter extends uniquely to a continuous homomorphism from the multiplicative group $K_\mathfrak{p}^\times$ of any finite extension $K_\mathfrak{p}$ of $\mathbb{Q}_p$ to the additive group of $K_\mathfrak{p}$, satisfying $\log_p(p) = 0$ and $\log_p(\zeta) = 0$ for every root of unity ζ; and $\exp_p$ is the formal inverse of $\log_p$ on its domain of convergence.

Strassmann's Theorem. The single most important analytic tool for our purposes is the following.

Theorem 9 (Strassmann's Theorem)**.** *Let $f(x) = \sum_{n\geq 0} c_n x^n$ be a non-zero p-adic power series converging on the closed unit disc $D(0,1)$. Let N be the largest index such that $|c_N|_p = \max_{n\geq 0} |c_n|_p$. Then f has at most N zeros in $D(0,1)$.*

Strassmann's Theorem is the p-adic analogue of the basic complex-analytic fact that a non-zero analytic function on a compact disc has only finitely many zeros, and is what underlies the finiteness assertion of the Skolem–Mahler–Lech Theorem (Theorem 16).

A closely related result, of which Strassmann's Theorem is essentially a corollary, is the *Weierstrass Preparation Theorem* for p-adic power series: any f as above can be written as $f(x) = u(x) \cdot P(x)$, where P is a polynomial of degree N with all roots in $D(0,1)$, and $u(x)$ is a power series converging on $D(0,1)$ with no zeros there. This finer statement underlies the complexity results of Section 5.4.

Mahler series. A *Mahler series* is a formal series

$$f(x) \;=\; \sum_{k\geq 0} c_k \binom{x}{k},$$

where $\binom{x}{k}$ is the polynomial $\binom{x}{k} := \frac{x(x-1)\cdots(x-k+1)}{k!}$ and $c_k \in \mathbb{Q}_p$ are such that $|c_k|_p \to 0$ as $k \to \infty$. A classical theorem of K. Mahler asserts that f as above defines a continuous function $f : \mathbb{Z}_p \to \mathbb{Q}_p$ and that every such function admits a unique representation in this form. Furthermore, if $v_p(c_k) \geq k$, then f defines an analytic function on $\mathbb{Z}_p$. Mahler series provide a particularly convenient framework in which to manipulate p-adic interpolations of integer sequences, and play a central role in the algorithms of the papers [13, 12] that we discuss in Section 5.4.

2.4 A Note on Models of Computation

Throughout the survey, all complexity statements are with respect to the standard Turing machine model, with inputs encoded in binary. A rational LRS is encoded by the rational coefficients of its defining recurrence and its rational initial values. The bit-size of an LRS is the total bit-length of this encoding.

The complexity classes we shall work with include the standard classes P, NP, coNP, PSPACE, the randomized classes RP and coRP, and, in Section 7, the relativized classes $\mathsf{NP}^{\mathsf{PosSLP}}$ and $\mathsf{coNP}^{\mathsf{PosSLP}}$. We shall also briefly mention the *Counting Hierarchy* CH, a hierarchy of complexity classes built from PP in the same way that the polynomial hierarchy is built from NP, which arises naturally in the analysis of the Positivity Problem at low orders (Sections 7.2 and 7.3).

A recurring technical issue, both for theoretical complexity and for practical implementation, is the manipulation of large numbers represented by straight-line programs or arithmetic circuits. The problem PosSLP asks whether a given arithmetic circuit over the integers, with operations $+$, $-$, $\times$, computes a positive integer; whether $\mathsf{PosSLP} \in \mathsf{P}$ is a major open problem [8]. The closely related decision problem EqSLP, which asks whether such a circuit evaluates to zero, is classically known to be in coRP [113].

3 Linear Recurrence Sequences

In this section we develop the basic theory of linear recurrence sequences. We state the results for rational LRS, which is our principal setting; the theory extends to LRS over any number field (Section 3.5) and, in its structural aspects, to arbitrary fields of characteristic 0. Most of the material is classical and can be found, for example, in [52, 118, 66]. Our purpose is

to fix the notation and establish the structural results—the exponential-polynomial form, the matrix form, the decomposition into non-degenerate parts, the basic closure properties, and the Skolem–Mahler–Lech Theorem—to which we shall appeal repeatedly in the remainder of the survey.

3.1 The Exponential-Polynomial Form

Recall from Section 1 that a rational LRS of order d is a sequence $\boldsymbol{u} = \langle u_n \rangle_{n=0}^{\infty}$ of rational numbers satisfying a recurrence

$$u_{n+d} \;=\; a_{d-1}\, u_{n+d-1} \;+\; \cdots \;+\; a_0\, u_n \qquad (n \in \mathbb{N}) \tag{3}$$

with $a_0, \ldots, a_{d-1} \in \mathbb{Q}$ and $a_0 \neq 0$, where d is least with this property, and that its characteristic polynomial is

$$p(x) \;=\; x^d - a_{d-1}\, x^{d-1} - \cdots - a_1\, x - a_0.$$

The roots of $p(x)$ in $\overline{\mathbb{Q}}$ are the *characteristic roots* of the LRS; the assumption $a_0 \neq 0$ ensures that none of these roots is zero. Throughout the survey, we denote the distinct characteristic roots by $\lambda_1, \ldots, \lambda_m$, and their respective multiplicities in $p(x)$ by $e_1, \ldots, e_m$, so that $\sum_{j=1}^{m} e_j = d$.

The following structural result expresses every LRS in closed form.

Theorem 10. *Let $\boldsymbol{u}$ be a rational LRS of order d, with distinct characteristic roots $\lambda_1, \ldots, \lambda_m$ of multiplicities $e_1, \ldots, e_m$. There exist unique polynomials $Q_1, \ldots, Q_m \in \overline{\mathbb{Q}}[x]$, with $\deg Q_j \leq e_j - 1$ for each j, such that*

$$u_n \;=\; \sum_{j=1}^{m} Q_j(n)\, \lambda_j^n \qquad (n \in \mathbb{N}). \tag{4}$$

The polynomials Q_j have coefficients in the splitting field of $p(x)$ over $\mathbb{Q}$, and are computable from $\boldsymbol{u}$ by solving a linear system.

Sketch. The space V of all sequences of complex numbers satisfying the recurrence (3) is a $\mathbb{C}$-vector space of dimension d, since a sequence in V is uniquely determined by its first d terms. For each $j \in \{1, \ldots, m\}$ and each integer $\ell \in \{0, 1, \ldots, e_j - 1\}$, the sequence $\langle n^\ell \lambda_j^n \rangle_{n=0}^{\infty}$ belongs to V; this is verified by substituting into the recurrence and applying the elementary identity $\sum_{k=0}^{d} \binom{n+k}{\ell} c_k \lambda^{n+k} = 0$ whenever λ is a root of multiplicity at least $\ell + 1$ of the polynomial $\sum_{k=0}^{d} c_k x^k$. The total number of such sequences is $\sum_{j=1}^{m} e_j = d$, and a Vandermonde-style argument shows that they are linearly independent over $\mathbb{C}$, so they form a basis of V. The claim follows by expressing $\boldsymbol{u}$ in this basis. The fact that the coefficients of the Q_j lie in the splitting field of $p(x)$ follows from the linearity of the system of equations expressing $u_0, \ldots, u_{d-1}$ in terms of those coefficients, and from the fact that this system has rational coefficients in any basis fixed by the Galois group. □

Since d is the order of $\boldsymbol{u}$, minimality forces each Q_j to be non-zero of degree exactly $e_j - 1$: otherwise $\boldsymbol{u}$ would satisfy a shorter recurrence.

The representation (4) is called the *exponential-polynomial form* of $\boldsymbol{u}$. Conversely, given any choice of distinct non-zero algebraic numbers $\lambda_1, \ldots, \lambda_m$ and polynomials $Q_1, \ldots, Q_m \in \overline{\mathbb{Q}}[x]$, the sequence defined by (4) is an LRS, with order at most $\sum_j (\deg Q_j + 1)$. From this perspective, an LRS may be specified either by its defining recurrence and initial values, or by its characteristic roots and the polynomials Q_j. Since converting between these representations chiefly involves polynomial root finding and algebraic-number manipulations (all of which can be performed in polynomial time [21]), either specification can be obtained from the other in polynomial time, and we shall pass freely between them.

An important notion regarding characteristic roots is that of *dominance*: a characteristic root λ is *dominant* if it has maximal modulus among all of the characteristic roots, i.e., if

$|\lambda| = \max_i |\lambda_i|$. Likewise, for $\mathfrak{p}$ a prime ideal, λ is $\mathfrak{p}$-*dominant* if it has maximal $\mathfrak{p}$-adic absolute value among all characteristic roots, i.e., if $|\lambda|_\mathfrak{p} = \max_i |\lambda_i|_\mathfrak{p}$.

Recall that an LRS is *simple* if its characteristic polynomial has no repeated roots, i.e., if $e_j = 1$ for every j. By Theorem 10, this is equivalent to requiring that each Q_j in (4) is a constant polynomial. A simple LRS thus admits a representation

$$u_n \;=\; \sum_{j=1}^{d} c_j\, \lambda_j^n \qquad (n \in \mathbb{N}) \tag{5}$$

with $c_j \in \overline{\mathbb{Q}}$ and the λ_j pairwise distinct.

Example 11. The Fibonacci sequence (Example 1) is simple, with $\lambda_1 = \varphi = (1+\sqrt{5})/2$, $\lambda_2 = \tilde{\varphi} = (1-\sqrt{5})/2$, and $c_1 = -c_2 = 1/\sqrt{5}$. The Lucas sequence $\langle 2, 1, 3, 4, 7, 11, 18, \ldots\rangle$, satisfying the same recurrence $L_{n+2} = L_{n+1} + L_n$ with initial values $L_0 = 2$, $L_1 = 1$, has characteristic roots φ and $\tilde{\varphi}$ (the same as the Fibonacci sequence) but coefficients $c_1 = c_2 = 1$, yielding the closed form $L_n = \varphi^n + \tilde{\varphi}^n$.

Example 12. Consider the order-3 LRS $\boldsymbol{u}$ defined by $u_{n+3} = 5u_{n+2} - 8u_{n+1} + 4u_n$, with initial values $u_0 = 0$, $u_1 = 1$, $u_2 = 5$. Its characteristic polynomial $x^3 - 5x^2 + 8x - 4 = (x-1)(x-2)^2$ has roots $\lambda_1 = 1$ (simple) and $\lambda_2 = 2$ (double); the LRS is therefore not simple. Its exponential-polynomial form is

$$u_n \;=\; 1 \cdot 1^n \;+\; (n-1) \cdot 2^n \qquad (n \in \mathbb{N}),$$

with $Q_1(n) = 1$ and $Q_2(n) = n - 1$. The non-constant polynomial Q_2, a hallmark of repeated roots, is responsible for several of the additional difficulties that arise in the analysis of non-simple LRS, as we shall see in Sections 5 and 7.

3.2 The Matrix Form

A useful alternative perspective on LRS is afforded by their realization as orbits of linear maps. Given an LRS $\boldsymbol{u}$ satisfying the recurrence (3), define the column vector $\boldsymbol{v}_n := (u_n, u_{n+1}, \ldots, u_{n+d-1})^\top$ and the *companion matrix* of the recurrence,

$$A \;:=\; \begin{pmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & & & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ a_0 & a_1 & a_2 & \cdots & a_{d-1} \end{pmatrix} \;\in\; \mathbb{Q}^{d\times d}. \tag{6}$$

Then $\boldsymbol{v}_{n+1} = A\boldsymbol{v}_n$ for all n, and consequently

$$u_n \;=\; \boldsymbol{e}_1^\top A^n\, \boldsymbol{v}_0 \qquad (n \in \mathbb{N}), \tag{7}$$

where $\boldsymbol{e}_1 = (1, 0, \ldots, 0)^\top$ is the first standard basis vector. The characteristic polynomial of A, in the sense of linear algebra, coincides with the characteristic polynomial of $\boldsymbol{u}$; its roots, with multiplicities, are precisely the eigenvalues of A.

More generally, one can consider sequences of the form $u_n = \boldsymbol{\alpha}^\top A^n \boldsymbol{\beta}$, where A is an arbitrary $d \times d$ matrix over $\mathbb{Q}$ and $\boldsymbol{\alpha}, \boldsymbol{\beta}$ are arbitrary d-dimensional rational vectors. Such sequences are exactly the LRS of order at most d, since the Cayley–Hamilton Theorem ensures that A^n satisfies a recurrence of order d, and conversely every LRS of order d is realized by the companion-matrix construction above. Thus the matrix form (7) is fully equivalent to the recurrence-and-initial-values form, modulo a polynomial-time change of representation.

The matrix perspective therefore leads naturally to the following alternative formulation of the Skolem Problem: given a square rational matrix A and rational vectors $\boldsymbol{\alpha}, \boldsymbol{\beta}$, does there

exist $n \in \mathbb{N}$ such that $\boldsymbol{\alpha}^\top A^n \boldsymbol{\beta} = 0$? In the rest of the survey, we shall use the two formulations interchangeably. The matrix view is also the natural starting point for the reachability and invariant-generation questions concerning linear dynamical systems that we discuss in Sections 4.4 and 9.

3.3 Non-Degeneracy and the Skolem–Mahler–Lech Theorem

We have already encountered the notion of non-degeneracy in the introduction; we now treat it more carefully.

Definition 13. An LRS $\boldsymbol{u}$ with distinct characteristic roots $\lambda_1, \ldots, \lambda_m$ is *non-degenerate* if no quotient λ_i/λ_j (for $i \neq j$) is a root of unity.

The relevance of this notion to the Skolem and Positivity Problems is best appreciated through the following construction. Suppose $\boldsymbol{u}$ is degenerate, so that $\lambda_i/\lambda_j = \zeta$ is a root of unity for some $i \neq j$, of order L say. Then for each residue class $r \in \{0, 1, \ldots, L-1\}$, the subsequence $\langle u_{Ln+r} \rangle_{n=0}^{\infty}$ is again an LRS, with characteristic roots among the λ_j^L; in this subsequence, the roots λ_i^L and λ_j^L coincide. Applying this construction with L a suitable common multiple of the orders of all root-of-unity quotients yields a decomposition of $\boldsymbol{u}$ into L non-degenerate subsequences indexed by residue class modulo L. The integer L is bounded effectively in terms of the degree of the characteristic roots; a classical theorem of A. Schinzel [109] guarantees that L may be taken at most $\exp\big(O(d \log d)\big)$, where d is the order of the original LRS. We summarize:

Proposition 14. *Every LRS $\boldsymbol{u}$ admits an effective decomposition into a finite interleaving of non-degenerate subsequences of the same or smaller order. Some of the subsequences may be identically zero. The number of subsequences in the decomposition is bounded by $\exp\big(O(d \log d)\big)$, and the decomposition itself is computable in time polynomial in the size of the resulting subsequences.*

Example 15. The LRS $\boldsymbol{u}$ given by $u_{n+2} = -u_n$ with $u_0 = 1$, $u_1 = 0$ is degenerate: its characteristic roots are $\pm i$, whose quotient -1 is a root of unity of order 2. The sequence is $\langle 1, 0, -1, 0, 1, 0, -1, 0, \ldots \rangle$, and decomposes into two interleaved non-degenerate subsequences: $\langle 1, -1, 1, -1, \ldots \rangle$ at even indices, and the zero sequence at odd indices. The Skolem Problem for the original LRS is then immediately resolved by inspecting the two subsequences.

The zero set of an arbitrary LRS admits the following classical description.

Theorem 16 (Skolem–Mahler–Lech)**.** *Let $\boldsymbol{u}$ be an LRS over a field of characteristic 0. Then the zero set*

$$Z(\boldsymbol{u}) \;:=\; \{n \in \mathbb{N} : u_n = 0\}$$

is the union of finitely many full arithmetic progressions together with a finite set. Equivalently, in any non-degenerate decomposition of $\boldsymbol{u}$, each non-zero non-degenerate subsequence has only finitely many zeros.

The Skolem–Mahler–Lech Theorem was proved by T. Skolem [115] for LRS over $\mathbb{Q}$, by K. Mahler [84] for algebraic LRS, and by C. Lech [73] in the generality stated above. All three proofs proceed via Skolem's p-adic method, which we develop in Section 5.2. The restriction to characteristic 0 is essential: over the field $\mathbb{F}_p(x)$, for instance, the order-3 LRS $u_n := (x+1)^n - x^n - 1$ vanishes precisely when n is a power of p, so that its zero set is neither finite nor a finite union of arithmetic progressions. The structure of zero sets in characteristic p is described by Derksen's theorem [49] in terms of p-automatic sets; see Section 9 for a brief discussion.

Proposition 14 and the Skolem–Mahler–Lech Theorem together account for the non-degeneracy hypothesis that pervades subsequent sections. By Proposition 14, deciding whether an arbitrary LRS has a zero reduces to deciding the same question on each of finitely many non-degenerate subsequences; and by the Skolem–Mahler–Lech Theorem, each such subsequence has only finitely

many zeros. The Skolem Problem is therefore equivalent to the problem of *computing* the finitely many zeros of a non-degenerate LRS, and we shall freely make use of this reduction whenever convenient.

A key effective refinement of the Skolem–Mahler–Lech Theorem concerns the *number* of zeros of a non-degenerate LRS. The following bound, successively sharpened by W. M. Schmidt [112], P. B. Allen [7], and F. Amoroso and E. Viada [10], will be invoked at several points; the explicit form below follows the presentation in [83].

Theorem 17. *There is an effectively computable function $N : \mathbb{N} \to \mathbb{N}$ such that any non-zero non-degenerate LRS of order d has at most $N(d)$ zeros. Explicitly, one may take $N(d) = \left(8d^{d}\right)^{8d^{6d}}$.*

The bound $N(d)$ depends only on the order of the LRS, and is valid for non-degenerate LRS over any field of characteristic 0. It is, on the other hand, certainly very far from sharp: already at order 3, the largest known number of zeros is six, attained by Berstel's sequence $u_{n+3} = 2u_{n+2} - 4u_{n+1} + 4u_n$ with initial values $0, 0, 1$, whose zero set is $\{0, 1, 4, 6, 13, 52\}$ (see [52]), whereas $N(3) = 216^{8 \cdot 27^{6}}$ exceeds $10^{10^{9}}$. Unfortunately—and this is precisely the content of the Skolem Problem—no effective bound is currently known on the *magnitude* of the zeros, and Theorem 17 does not directly help with their computation: even if one knows that there are at most N zeros, one does not know in what range to search for them.

3.4 Closure Properties

The class of LRS over a fixed field is closed under several natural operations.

Proposition 18. *Let $\boldsymbol{u}$ and $\boldsymbol{v}$ be LRS of orders d_u and d_v respectively. Then:*

1. *the termwise sum $\boldsymbol{u} + \boldsymbol{v} := \langle u_n + v_n \rangle_{n=0}^{\infty}$ is an LRS of order at most $d_u + d_v$,*
2. *the Hadamard product $\boldsymbol{u} \odot \boldsymbol{v} := \langle u_n v_n \rangle_{n=0}^{\infty}$ is an LRS of order at most $d_u \cdot d_v$,*
3. *for any $a \in \mathbb{Q}$, the scalar multiple $\langle a u_n \rangle_{n=0}^{\infty}$ is an LRS of order at most d_u,*
4. *for any $a, k \in \mathbb{N}$, the shifted subsequence $\langle u_{an+k} \rangle_{n=0}^{\infty}$ is an LRS of order at most d_u.*

Over the field of rational numbers, all of the corresponding constructions are effective and computable in polynomial time.

The proofs are direct consequences of the exponential-polynomial representation: if $\boldsymbol{u}$ has characteristic roots $\{\lambda_i\}$ and $\boldsymbol{v}$ has characteristic roots $\{\mu_j\}$, then $\boldsymbol{u} + \boldsymbol{v}$ has characteristic roots a subset of $\{\lambda_i\} \cup \{\mu_j\}$, and $\boldsymbol{u} \odot \boldsymbol{v}$ has characteristic roots a subset of $\{\lambda_i \mu_j\}$. The order bounds are sharp in general, though not always attained.

A consequence used several times below is that, given an LRS $\boldsymbol{u}$, the sequence $\langle u_n^2 \rangle_{n=0}^{\infty} = \boldsymbol{u} \odot \boldsymbol{u}$ is also an LRS, of order at most $\binom{d_u+1}{2}$. This is the construction underlying the Skolem-to-Positivity reduction sketched in Section 1, and it will reappear in Section 7.

3.5 Algebraic LRS

While most of our results concern integer or rational LRS, we shall occasionally consider *algebraic LRS*: LRS whose recurrence coefficients and initial values lie in a fixed number field K. The exponential-polynomial form, the matrix form, the non-degenerate decomposition, and the Skolem–Mahler–Lech Theorem all carry over without change to this setting; only the complexity analysis is affected, and even then in a controlled way (the bit-size of K enters as an additional parameter). P. Bacik's recent result [11] for the Skolem Problem at order 4 over the algebraic numbers, which we discuss in Section 5.3, is naturally formulated in this setting.

4 Certificates of Zero-Freeness and Non-Negativity

Before plunging into the technical heart of the survey, we pause to introduce what we regard as a unifying conceptual framework: the notion of a *certificate* for the Skolem and Positivity Problems. The framework is informal, but it organizes both classical results and contemporary developments into three recognizable patterns, and makes it easier to articulate where existing techniques succeed, where they fail, and what kind of new technique would be required to push the state of the art further.

4.1 The Certificate Framework

At a high level, deciding the Skolem Problem on a given input $\boldsymbol{u}$ amounts to one of two outcomes: either an integer n with $u_n = 0$ is exhibited (a *positive witness*), or the non-existence of such an integer is established by some finite mathematical argument (a *negative certificate*). Positive witnesses are unproblematic: an integer n together with the verification $u_n = 0$ is a finite, effectively checkable object. The substantive content of the Skolem Problem lies in the negative direction: how does one establish that $\boldsymbol{u}$ has no zero, given that there are infinitely many indices to check? Each unconditional decidability result in Section 5, and each conditional decidability result in Section 6, can be read as a procedure for producing such a negative certificate (or, more precisely, for deciding which of the two outcomes obtains). Three broad classes of certificates have been identified to date:

1. *Separation-bound certificates*, which establish $|u_n| > 0$ for all n beyond an effectively computable threshold N, by exhibiting a quantitative lower bound on $|u_n|$ that holds for all such n. This lower bound is typically obtained either from the Archimedean structure of the LRS (via Baker's Theorem on linear forms in logarithms) or from its p-adic analogue.

2. *Modular invariants*, which establish $u_n \neq 0$ for *all* $n \in \mathbb{N}$ (or $\mathbb{Z}$, in the case of an LRBS) by exhibiting a modulus $m \geq 2$ such that the periodic sequence $\langle u_n \bmod m \rangle$ is everywhere non-zero. The witnessing modulus m is the certificate; checking it involves a single period of $\boldsymbol{u}$ modulo m, which is a finite computation.

3. *Semialgebraic invariants*, which establish $u_n \neq 0$ by exhibiting, in the matricial formulation of the Skolem Problem (Section 3.2), an A-invariant semialgebraic set $I \subseteq \mathbb{R}^d$ that contains the initial point $\boldsymbol{v}_0$ and is disjoint from the hyperplane $\{\boldsymbol{x} : \boldsymbol{e}_1^\top \boldsymbol{x} = 0\}$, which the orbit visits precisely at the zeros of $\boldsymbol{u}$.

Certification applies equally to the Positivity Problem, with non-negativity in place of non-zeroness—though with an important caveat, to which we return in Section 4.5.

We now describe each class in turn.

4.2 Separation-Bound Certificates

The simplest and historically earliest class of certificates relies on the following observation: if one can establish, for a given LRS $\boldsymbol{u}$, an effective bound $N \in \mathbb{N}$ such that $|u_n| > 0$ for all $n \geq N$, then deciding whether $\boldsymbol{u}$ has a zero reduces to evaluating $u_0, u_1, \ldots, u_{N-1}$ and checking each. Such zero-separation bounds come in two kinds, resting either on the *Archimedean (Euclidean) absolute value* or on one of its *non-Archimedean (p-adic)* counterparts. Both kinds are deployed in the low-order decidability results of Section 5.3, and both are illustrated concretely in Section 8.2.

A disadvantage of both Archimedean and non-Archimedean separation bounds is that, in the current state of knowledge, their effective application typically requires strong structural restrictions on the dominant roots, as in the MSTV class (Section 5.3). For general LRS at high orders, such separation bounds remain elusive.

4.3 Modular Invariants

A second class of certificates of zero-freeness proceeds through modular arithmetic. Suppose $\boldsymbol{u}$ is an integer LRS, and suppose one can exhibit a modulus $m \geq 2$ such that, viewed in $\mathbb{Z}/m\mathbb{Z}$, the sequence $\boldsymbol{u}$ takes only non-zero values. The reduction modulo m is ultimately periodic—indeed purely periodic when $\gcd(a_0, m) = 1$, the companion matrix then being invertible modulo m, with period at most the order of that matrix in $\mathrm{GL}_d(\mathbb{Z}/m\mathbb{Z})$—so checking that it is everywhere non-zero is a finite computation. If the check succeeds, the conclusion is that $\boldsymbol{u}$ has no integer zero, since any integer n with $u_n = 0$ would give $u_n \equiv 0 \pmod{m}$. This furnishes a *modular-invariant* certificate: the witness is the modulus m together with the verification that the reduction contains no zero. (It would in fact suffice for the reduction to be non-zero from some index onward, the finitely many earlier indices being checked directly.) Though superficially quite different from separation bounds, the two classes are related: a witnessing modulus of the form p^k amounts precisely to the p-adic separation bound $|u_n|_p > p^{-k}$, valid for all $n \in \mathbb{N}$.

When modular invariants suffice. The Exponential Local-Global Principle (Section 6.1) asserts, for simple LRBS, that the modular-invariant approach is in fact *always* sufficient: every zero-free simple LRBS admits a witnessing modulus m. The principle is conjectural in general, but is known unconditionally in a number of special cases (orders 2 and certain order-3 families).

When modular invariants fail. For non-simple LRS, modular invariants can fail to exist even when the LRS is zero-free: the LRBS $u_n = (2n+1)\,2^n$, for instance, is everywhere non-zero on $\mathbb{Z}$, yet has zeros modulo every admissible modulus (see Example 28 in Section 6.1 for the details). A more elaborate and algorithmically pertinent example is the order-6 LRS

$$u_n \;=\; 2(-4+7i)^n + 2(-4-7i)^n + 4(8+i)^n + 4(8-i)^n + n, \tag{8}$$

which we shall return to in Section 8.3: this LRS possesses zeros modulo every $m \geq 2$ (Proposition 53), and in particular admits no modular invariant certifying its zero-freeness. Whether the LRS is genuinely zero-free on $\mathbb{N}$ is open at the time of writing, and is the kind of question for which a combination of certificate techniques (or potentially new techniques altogether) may be required.

4.4 Semialgebraic Invariants

A third and structurally distinct class of certificates proceeds via the geometric viewpoint on the Skolem Problem developed in Section 3.2. Recall that an LRS $\boldsymbol{u}$ of order d may be realized in matrix form as $u_n = \boldsymbol{e}_1^\top A^n \boldsymbol{v}_0$, where A is the companion matrix of $\boldsymbol{u}$ and $\boldsymbol{v}_0 \in \mathbb{Q}^d$ is the initial state. The Skolem Problem then becomes the question of whether the orbit $\{A^n\boldsymbol{v}_0 : n \in \mathbb{N}\}$ ever intersects the hyperplane $H := \{\boldsymbol{x} \in \mathbb{R}^d : \boldsymbol{e}_1^\top \boldsymbol{x} = 0\}$. A *semialgebraic invariant* for this question is a semialgebraic set $I \subseteq \mathbb{R}^d$ satisfying the following three conditions:

1. $\boldsymbol{v}_0 \in I$,
2. $AI \subseteq I$, i.e., I is closed under the dynamics, and
3. $I \cap H = \emptyset$.

The existence of such an I constitutes a finite, semialgebraic certificate of zero-freeness: by induction on n, $A^n\boldsymbol{v}_0 \in I$ for all $n \geq 0$, hence $A^n\boldsymbol{v}_0 \notin H$, hence $u_n \neq 0$.

Algebraic invariants as a special case. An important special case is that of *algebraic invariants*, in which I is taken to be an algebraic set. Here a fundamental result is that the *strongest* algebraic invariant (smallest algebraic set) containing a given initial state is effectively computable [59, 60]. Algebraic invariants are rarely sufficient on their own to establish zero-freeness of non-trivial LRS, but they form a useful subclass and play a role in several recent algorithmic advances; a worked example of an algebraic invariant certifying zero-freeness appears in Section 8.1.

Decidability of existence of semialgebraic invariants. Unfortunately, strongest semialgebraic invariants do not generally exist. A natural algorithmic question is thus whether, given an LRS $\boldsymbol{u}$ and a target hyperplane (or, in the Positivity case, a target half-space), one can decide whether *some* semialgebraic invariant of the kind above exists. The following result settles this in the affirmative [9].

Theorem 19. *There is an effective procedure that, given an LRS $\boldsymbol{u}$ in matrix form $(A, \boldsymbol{v}_0)$ and a hyperplane or half-space $H \subseteq \mathbb{R}^d$, decides whether there exists a semialgebraic invariant $I \subseteq \mathbb{R}^d$ satisfying $\boldsymbol{v}_0 \in I$, $AI \subseteq I$, and $I \cap H = \emptyset$.*

Theorem 19 not only provides a decision procedure for the existence of a separating semialgebraic invariant, but also guarantees that, in positive instances, such an invariant can be explicitly constructed. Note, however, that in negative instances, no conclusion can be drawn regarding the zero-freeness (or positivity) of the corresponding LRS. We shall return to this point in Section 8.

4.5 Certificates for Positivity

The Positivity Problem also admits classes of certificates, but with one important asymmetry: *modular and p-adic certificates are of no use for Positivity.* The reason is straightforward: modular reduction and p-adic absolute value are blind to the order on $\mathbb{R}$, and retain no information whatsoever about the signs of the u_n.

The certificate landscape for Positivity is therefore narrower than for Skolem. The two relevant classes are:

1. *Archimedean separation bounds.* The unconditional decidability results of [99] for general LRS at orders ≤ 5, and [98] for simple LRS at orders ≤ 9, all proceed by Archimedean separation bounds via Baker's Theorem: the dominant part of the sequence, whose sign governs that of u_n, is bounded away from zero and raced against the exponentially smaller non-dominant part.

2. *Semialgebraic invariants.* For Positivity, an A-invariant semialgebraic set $I \subseteq \mathbb{R}^d$ that is disjoint from the open half-space $\{\boldsymbol{x} : \boldsymbol{e}_1^\top \boldsymbol{x} < 0\}$ certifies $u_n \geq 0$ for all n in the same manner as for Skolem. Theorem 19 applies with the half-space in place of the hyperplane; the decidability of existence of a semialgebraic invariant for Positivity is therefore on the same footing as for Skolem.

This narrowness is, in our view, one of the chief structural reasons that Positivity is genuinely harder than Skolem. The Diophantine hardness barriers of Section 7.4, and the order-10 hard instance for simple LRS that we exhibit in Section 8, can be read as concrete witnesses to the limits of the available certification techniques.

4.6 Beyond Existing Certificate Classes

We close this section with two observations.

First, the three certificate classes we have identified are not mutually exclusive: a given LRS may admit certificates of more than one kind, and a successful decidability proof on a given class of inputs may proceed by combining several. The analysis of the Skolem Problem at order 3 and order 4, for example, combines Archimedean separation (via Baker) with p-adic separation (via p-adic Baker).

Second, the certificate classes are not exhaustive by any means: in Section 8 we exhibit 'hard' instances—the order-6 sequence (8) and an order-10 simple sequence—for which the existing certificate classes either provably fail or lie beyond current techniques, so that we are at present unable to establish zero-freeness and positivity respectively. Note that these instances are not merely isolated open problems. Rather, they serve a structural purpose, helping calibrate the limits of existing certificate classes, and a successful argument settling their Skolem or Positivity status may well introduce new certificate classes that apply more broadly than to the original instances alone. In this sense, the analysis of hard instances is a potential pathway to general decidability results.

5 The Skolem Problem: Unconditional Results

This section is devoted to the unconditional state of the art for the Skolem Problem, which we recall asks, given an LRS $\boldsymbol{u}$ over $\mathbb{Q}$ (or more generally over a number field), whether there exists $n \in \mathbb{N}$ such that $u_n = 0$. We trace the historical arc of the field and present the decidability and complexity results known to hold without recourse to any unproved number-theoretic conjecture; conditional results and alternative approaches are the subject of Section 6.

In the language of the certificate framework introduced in Section 4, every result of this section can be read as a procedure for producing a separation-bound certificate, of either Archimedean or non-Archimedean kind, on a specific structural class of LRS. The MSTV class of Section 5.3 captures the regime in which both flavors can be combined to greatest effect.

5.1 Statement, History, and Basic Equivalences

We have already met the Skolem Problem in Section 1; let us now examine its various equivalent formulations and their early history.

Equivalent formulations. The Skolem Problem admits several equivalent presentations, all of which we shall refer to interchangeably.

1. *The recurrence form.* Given a rational LRS $\boldsymbol{u}$ specified by its recurrence coefficients $a_0, \ldots, a_{d-1} \in \mathbb{Q}$ and initial values $u_0, \ldots, u_{d-1} \in \mathbb{Q}$, decide whether there exists $n \in \mathbb{N}$ with $u_n = 0$.

2. *The matrix form.* Given a square rational matrix $A \in \mathbb{Q}^{d\times d}$ and rational vectors $\boldsymbol{\alpha}, \boldsymbol{\beta} \in \mathbb{Q}^d$, decide whether there exists $n \in \mathbb{N}$ with $\boldsymbol{\alpha}^\top A^n \boldsymbol{\beta} = 0$.

3. *The hyperplane form.* Given a square rational matrix A, a rational vector $\boldsymbol{\beta}$, and a rational hyperplane $H \subseteq \mathbb{Q}^d$, decide whether there exists $n \in \mathbb{N}$ with $A^n \boldsymbol{\beta} \in H$.

4. *The exponential-polynomial form.* Given distinct non-zero algebraic numbers $\lambda_1, \ldots, \lambda_m$ and polynomials $Q_1(x), \ldots, Q_m(x)$, where $Q_j(x)$ has coefficients in $\mathbb{Q}(\lambda_j)$, decide whether there exists $n \in \mathbb{N}$ with $\sum_j Q_j(n)\lambda_j^n = 0$.

The equivalence of 1, 2, and 4 was established in Section 3; the equivalence with 3 is immediate, since $\boldsymbol{\alpha}^\top A^n \boldsymbol{\beta} = 0$ if and only if $A^n \boldsymbol{\beta}$ lies in the hyperplane $\{\boldsymbol{x} : \boldsymbol{\alpha}^\top \boldsymbol{x} = 0\}$. All four reductions are effective in polynomial time.

From algebraic to integer LRS. The Skolem Problem is also robust with respect to the field of definition. Zeros are preserved under scaling, so the rational and integer versions of the problem coincide (Section 1); more substantially, the algebraic case reduces to the integer one as well [80].

Proposition 20. *Let $\boldsymbol{v}$ be an LRS of order d over $\overline{\mathbb{Q}}$, defined over a number field K of degree D. One can effectively construct an integer LRS $\boldsymbol{u}$ of order at most d^D with exactly the same zero set as $\boldsymbol{v}$; moreover, if $\boldsymbol{v}$ is simple, then so is $\boldsymbol{u}$. Consequently, the Skolem Problem for algebraic LRS is many-one reducible to the Skolem Problem for integer LRS, and likewise the Skolem Problem for simple algebraic LRS to the Skolem Problem for simple integer LRS.*

The construction is simply the field norm: writing $N_K(x) := \prod_\sigma \sigma(x)$, with the product ranging over the D field embeddings $\sigma : K \to \mathbb{C}$, one sets $u_n := N_K(v_n)$, scaling afterward to clear denominators, as in Section 1. Each $\langle \sigma(v_n) \rangle_{n=0}^{\infty}$ is again an LRS of order d, simple whenever $\boldsymbol{v}$ is, so the closure properties of Section 3.4 bound the order of $\boldsymbol{u}$ by d^D; and since field embeddings fix 0, the sequences $\boldsymbol{u}$ and $\boldsymbol{v}$ have precisely the same zeros.

A further equivalence, of independent interest, links the Skolem Problem to $\mathbb{Z}$-weighted automata [26, 114]. A $\mathbb{Z}$-weighted automaton is a finite automaton each of whose transitions carries an integer weight; the weight of a path is the product of the weights of its transitions, and the weight of a word is the sum of the weights of all accepting paths labeled by that word. A classical theorem of M. P. Schützenberger asserts that the function $w \mapsto \mathrm{weight}(w)$ from Σ^* to $\mathbb{Z}$ is realized by a $\mathbb{Z}$-weighted automaton if and only if it is a *rational formal series*, and the restriction of such a series to a unary alphabet $\Sigma = \{a\}$ yields, after identification of a^n with $n \in \mathbb{N}$, precisely an integer LRS. The Skolem Problem is therefore equivalent to deciding, for a given unary $\mathbb{Z}$-weighted automaton, whether some word has weight zero.

A natural restriction of the Skolem Problem to a finite range yields the following.

Problem 21 (The Bounded Skolem Problem)**.** *Given an LRS $\boldsymbol{u}$ and a positive integer N (encoded in binary), does there exist $n \in \{0, 1, \ldots, N\}$ such that $u_n = 0$?*

The Bounded Skolem Problem is of course decidable, since the search space is finite. However, as N is encoded in binary, a naive term-by-term examination requires exponential time. We shall return to these matters when we discuss recent complexity results in Section 5.4.

Hardness. On the lower-bound side, the Skolem Problem is known to be NP-hard, by a reduction due to V. D. Blondel and N. Portier [34] from the (complement of the) universality problem for unary NFAs (itself coNP-complete, see [119]). The reduction is short and informative; we sketch its main idea here. Let A be the 0-1 adjacency matrix of a unary NFA (over singleton alphabet $\{a\}$, say), and let $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ be the vectors of initial and accepting states, respectively. For any $n \in \mathbb{N}$, the expression $\boldsymbol{\alpha}^\top A^n \boldsymbol{\beta}$ counts the number of accepting paths of length n from an initial state to an accepting state, and the word a^n is accepted by the NFA precisely if this number is strictly positive. Thus the NFA fails to be universal if and only if there is some n for which $\boldsymbol{\alpha}^\top A^n \boldsymbol{\beta} = 0$. The required conclusion immediately follows.

Note that an NFA is not universal if and only if it rejects some word of at most exponential length, as the determinization procedure via subset construction readily shows. It follows that the above reduction adapts in a straightforward way to the Bounded Skolem Problem, yielding NP-hardness for that variant as well. No hardness result stronger than NP-hardness is currently known for the Skolem Problem.

Alternative proofs of NP-hardness can be found in [4, 3]; in particular, [4] establishes NP-completeness for the Skolem Problem in which all characteristic roots are complex roots of unity.

Skolem's 1934 paper. The Skolem Problem traces back to a 1934 paper of T. Skolem [115], in which he showed that any non-degenerate integer LRS has only finitely many zeros. Skolem's proof introduced what is now known as the *p-adic method*, the central technique to which we shall return repeatedly. Skolem himself did not pose the problem of *computing* the zeros, but the question is implicit in his work; it was first stated explicitly in print, to the best of our knowledge, by several authors in the mid-to-late 1970s, such as M. Mignotte [90], A. Salomaa, and M. Soittola [75, 108], and has been a celebrated open problem ever since. (The name of C. Pisot is occasionally associated with the Skolem Problem, but despite our best efforts we have been unable to find reliable sources in the literature concretely linking Pisot to the study of the Skolem Problem.)

5.2 The Skolem–Mahler–Lech Theorem and Skolem's p-adic Method

Recall the Skolem–Mahler–Lech Theorem (Theorem 16): the zero set of an LRS over a field of characteristic 0 is a union of finitely many full arithmetic progressions together with a finite set. We now sketch Skolem's original argument, pertaining to integer LRS.

Proof. Let $\boldsymbol{u}$ be an integer LRS. Write $\boldsymbol{u}$ in matrix form $u_n = \boldsymbol{e}_1^\top A^n \boldsymbol{v}_0$, where A is the companion matrix of $\boldsymbol{u}$ and $\boldsymbol{v}_0 \in \mathbb{Z}^d$ is the vector of initial terms. Then $\det(A) = \pm a_0$ is non-zero.

Let $p \in \mathbb{N}$ be a prime that does not divide $\det(A)$. Then A is invertible modulo p and there exists a positive integer L such that $A^L \equiv I \bmod p$, that is, $A^L = I + pB$ for some integer matrix $B \in \mathbb{Z}^{d\times d}$. Then, for $r \in \{0, \ldots, L-1\}$ we have:

$$\begin{aligned} u_{Ln+r} &= \boldsymbol{e}_1^\top A^{Ln+r} \boldsymbol{v}_0 \\ &= \boldsymbol{e}_1^\top A^r (A^L)^n \boldsymbol{v}_0 \\ &= \boldsymbol{e}_1^\top A^r (I + pB)^n \boldsymbol{v}_0 \\ &= \sum_{k=0}^{n} \binom{n}{k} p^k \boldsymbol{e}_1^\top A^r B^k \boldsymbol{v}_0. \end{aligned}$$

Writing $c_{r,k} := p^k \boldsymbol{e}_1^\top A^r B^k \boldsymbol{v}_0$, we have that $v_p(c_{r,k}) \geq k$. Hence the Mahler series expansion

$$f_r(x) := \sum_{k=0}^{\infty} c_{r,k} \binom{x}{k}$$

defines a p-adic analytic function $f_r : \mathbb{Z}_p \to \mathbb{Z}_p$ (see Section 2.3). By definition of f_r we have $f_r(n) = u_{Ln+r}$ for all $n \in \mathbb{N}$. If f_r is identically zero, this accounts for one of the full arithmetic progressions of zeros. Otherwise, by Strassmann's Theorem (Theorem 9), we have that f_r has finitely many zeros on $\mathbb{Z}_p$. Hence the zero set of our original LRS $\boldsymbol{u}$ consists of finitely many full arithmetic progressions together with a finite set, as required. □

Effectivity, and its absence. The proof above is ineffective in a precise sense: while Strassmann's Theorem provides an effective upper bound on the number of zeros of each f_r in $\mathbb{Z}_p$, it does not provide any information on *where* those zeros lie, and in particular, for any of these zeros, whether it belongs to $\mathbb{Z}$ (or $\mathbb{N}$). The state-of-the-art unconditional algorithms to which we turn shortly sidestep this difficulty by exploiting Baker's Theorem and a careful case analysis on the structure of the characteristic roots, but they apply only at low orders.

Effective bounds on the number of zeros. While effective bounds on the *magnitude* of the zeros of a non-degenerate LRS remain elusive, effective bounds on their *number* have been known for some time. We have already mentioned Theorem 17, which gives a bound doubly exponential in d, due ultimately to F. Amoroso and E. Viada [10]. Through the notion of Universal Skolem Sets (Section 6.4), we shall see how bounds on the number of zeros can be transformed into bounds on their magnitude under certain arithmetic assumptions.

The role of characteristic 0. The restriction to characteristic 0 in the Skolem–Mahler–Lech Theorem is essential: in characteristic p, the zero set of an LRS need not be a union of a finite set and finitely many arithmetic progressions, as the example following Theorem 16 shows. The structure of zero sets in positive characteristic is nonetheless well understood, and the Skolem Problem is decidable in that setting, thanks to an effective structure theorem of H. Derksen [49]; see Section 9.

5.3 Decidability at Low Orders and the MSTV Class

We now turn to the unconditional decidability results for the Skolem Problem. Despite very recent progress, the picture remains rather limited: decidability is known essentially only for LRS of order at most 4.

Easy cases. For LRS of order 0 or 1, sequences are of the form $u_n = c\lambda^n$ for some constants c, λ, and a zero exists if and only if $c = 0$, in which case the sequence is identically zero. For LRS of order 2, the hardest subcase involves sequences having closed form $u_n = c_1\lambda_1^n + c_2\lambda_2^n$, with all constants non-zero algebraic numbers. In turn, finding n such that $u_n = 0$ involves solving the exponential Diophantine equation $(\lambda_1/\lambda_2)^n = -c_2/c_1$, an elementary exercise in algebraic number theory (the procedure can moreover be carried out in polynomial time) [43].

The first substantial results appear at orders 3 and 4, and were established in the mid-1980s by M. Mignotte, T. N. Shorey, and R. Tijdeman [91] and independently by N. K. Vereshchagin [125]. Rather than retracing the original case-by-case treatments, we present the underlying technical result in the form in which it is understood today: as a decidability theorem for a class of LRS singled out by the structure of their dominant characteristic roots, nowadays called the *MSTV class* [76] after the four aforementioned authors. Our presentation follows the recent self-contained exposition of Y. Bilu [30], to which we refer the reader for complete proofs.

The MSTV class. Let $\boldsymbol{u}$ be a non-degenerate LRS with distinct characteristic roots $\lambda_1, \ldots, \lambda_m$, and let K be a number field containing all of them. Recall from Section 2.1 that the places of K comprise the Archimedean places, arising from the embeddings of K into $\mathbb{C}$, and the non-Archimedean places, arising from the non-zero prime ideals of $\mathcal{O}_K$. Generalizing the terminology of Section 3.1, given a place v of K with associated absolute value $|\cdot|_v$, we say that a characteristic root λ_j is *v-dominant* if $|\lambda_j|_v = \max_i |\lambda_i|_v$.

Definition 22. Let $\boldsymbol{u}$ be a non-degenerate LRS over $\overline{\mathbb{Q}}$ whose characteristic roots lie in the number field K. We say that $\boldsymbol{u}$ belongs to the *MSTV class* if at least one of the following two conditions holds:

1. there exists an Archimedean place v of K such that $\boldsymbol{u}$ has at most three v-dominant characteristic roots;
2. there exists a non-Archimedean place v of K such that $\boldsymbol{u}$ has at most two v-dominant characteristic roots.

Since absolute values extend and restrict along field extensions in a manner that preserves the induced ordering of the $|\lambda_j|_v$, membership in the MSTV class does not depend on the choice of the number field K. Moreover, when the characteristic polynomial has rational coefficients—in particular, for LRS over $\mathbb{Q}$—every embedding of K into $\mathbb{C}$ permutes the characteristic roots, and the first condition then simply asks that at most three of the roots be dominant in modulus, in the sense of Section 3.1.

Theorem 23 (MSTV Theorem)**.** *Let $\boldsymbol{u}$ be a non-degenerate LRS in the MSTV class. Then one can effectively compute all zeros of $\boldsymbol{u}$. In particular, the Skolem Problem is decidable on the MSTV class.*

Two remarks are in order. First, the theorem delivers more than a yes/no answer to the Skolem Problem: it places an effective bound on the magnitude of any zero, whence the full (finite) zero set can be computed by direct search. Second, the MSTV class contains LRS of arbitrarily high order, provided their dominant-root structure is sufficiently constrained; the restriction to low orders in the results below enters only through the task of exhibiting a suitable place v.

The role of Baker's Theorem. Let us sketch how the two versions of Baker's Theorem from Section 2.2 combine to prove Theorem 23. Fix a place v as in Definition 22, and renumber the characteristic roots so that $|\lambda_1|_v \geq |\lambda_2|_v \geq \cdots \geq |\lambda_m|_v$. The v-dominant roots are then $\lambda_1, \ldots, \lambda_k$, where $k \leq 3$ if v is Archimedean, and $k \leq 2$ if v is non-Archimedean. Splitting the exponential-polynomial representation of $\boldsymbol{u}$ accordingly, write

$$u_n \;=\; A(n) + B(n), \qquad A(n) := \sum_{j=1}^{k} Q_j(n)\,\lambda_j^n, \qquad B(n) := \sum_{j=k+1}^{m} Q_j(n)\,\lambda_j^n.$$

(Throughout the sketch we tacitly set aside the finitely many indices n at which some $Q_j(n)$ vanishes; these are effectively computable and are checked directly.) The tail $B(n)$ is exponentially smaller than $|\lambda_1|_v^n$: elementary estimates yield effective constants C_1, D such that $|B(n)|_v \leq C_1\, n^D\, |\lambda_{k+1}|_v^n$ for all $n \geq 2$, where $|\lambda_{k+1}|_v < |\lambda_1|_v$. The entire difficulty is to establish a matching effective lower bound of the shape

$$|A(n)|_v \;\geq\; |\lambda_1|_v^n\, e^{-C\,(\log n)^2} \tag{9}$$

for all sufficiently large n at which $A(n) \neq 0$, with the vanishing of $A(n)$ handled separately. Granted such a bound, we have $|A(n)|_v > |B(n)|_v$, and hence $u_n \neq 0$, for all n beyond an effectively computable threshold N; the finitely many remaining indices $n \leq N$ are then checked by direct evaluation. In the vocabulary of Section 4, the bound (9), together with the threshold N, constitutes a separation-bound certificate, of Archimedean or non-Archimedean type according to the nature of v.

The lower bound (9) is established as follows. When $k = 1$, the quantity $|A(n)|_v = |Q_1(n)|_v\, |\lambda_1|_v^n$ is bounded below directly: the height of the algebraic number $Q_1(n)$ grows only logarithmically in n, and Liouville's inequality (bounding a non-zero algebraic number from below in terms of its height and degree) yields $|Q_1(n)|_v \geq e^{-C \log n}$ whenever $Q_1(n) \neq 0$.

When $k = 2$, factoring out the first term gives

$$|A(n)|_v \;=\; |Q_1(n)|_v\, |\lambda_1|_v^n\, \big|\alpha_n \lambda^n - 1\big|_v, \qquad \text{where } \alpha_n := -\frac{Q_2(n)}{Q_1(n)}, \quad \lambda := \frac{\lambda_2}{\lambda_1}.$$

If $\alpha_n \lambda^n = 1$ then $n\, h(\lambda) = h(\alpha_n) = O(\log n)$; since λ is not a root of unity (by non-degeneracy), its height is effectively bounded away from zero, and n is effectively bounded. Otherwise, Baker's Theorem—in the Archimedean form of Theorem 7 when v is Archimedean, and in the p-adic form of Theorem 8 when v is non-Archimedean—provides a lower bound $|\alpha_n \lambda^n - 1|_v \geq e^{-C(h(\alpha_n)+1)\log n}$, and since $h(\alpha_n) = O(\log n)$, the bound (9) follows.

When $k = 3$—so that v is Archimedean—we may regard K as a subfield of $\mathbb{C}$ and drop the subscript v. Two new issues arise. First, the dominant part $A(n)$ may vanish for certain n even with all $Q_j(n) \neq 0$; but any such n is effectively bounded: dividing the equation $A(n) = 0$ by its third term and evaluating at a place at which λ_2/λ_3 is strictly small (such a place exists, in an effective sense, because λ_2/λ_3 is not a root of unity), one obtains a two-term expression that decays exponentially in n at that place, while Baker's Theorem bounds it below by $e^{-C(\log n)^2}$; comparing the two bounds confines n to an effectively computable range. Second, for the main lower bound one writes

$$A(n) \;=\; Q_1(n)\,\lambda_1^n\,\big(z_1(n) - z_0(n)\big), \qquad z_0(n) := -\frac{Q_2(n)}{Q_1(n)}\left(\frac{\lambda_2}{\lambda_1}\right)^n, \quad z_1(n) := 1 + \frac{Q_3(n)}{Q_1(n)}\left(\frac{\lambda_3}{\lambda_1}\right)^n,$$

and observes that $z_0(n)$ lies on a circle centered at 0 and $z_1(n)$ on a circle centered at 1, the radii of the circles being governed by $|Q_2(n)/Q_1(n)|$ and $|Q_3(n)/Q_1(n)|$ respectively. If the two circles are well separated, then $|z_1(n) - z_0(n)|$ is trivially bounded below. Otherwise—and this is a geometric observation going back to F. Beukers and R. Tijdeman [29]—the two points can be close to one another only when both are close to an intersection point ς of the two circles, and each of the distances $|z_0(n) - \varsigma|$ and $|z_1(n) - \varsigma|$ is again of the two-term shape $|\alpha\lambda^n - 1|$, amenable to Baker's Theorem exactly as in the case $k = 2$. We refer to [30] for the details.

Low orders lie in the MSTV class. An LRS with at most three distinct characteristic roots trivially belongs to the MSTV class. The Skolem Problem is therefore decidable for non-degenerate LRS of order at most 3 over the algebraic numbers—and hence, via the effective decomposition of Proposition 14, for arbitrary algebraic LRS of order at most 3.

The case of four distinct characteristic roots is genuinely harder: an LRS may have four dominant roots at every Archimedean place, ruling out the first condition of Definition 22, so that one must go looking for a suitable non-Archimedean place. This analysis was carried out for *real-valued* algebraic LRS of order at most 4 in the original works of the mid-1980s; the general complex-algebraic case was settled only recently [11], by means of the following lemma.

Lemma 24. *Every non-degenerate LRS over $\overline{\mathbb{Q}}$ with four distinct characteristic roots—in particular, every non-degenerate simple algebraic LRS of order 4—belongs to the MSTV class. Concretely, if $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \in K^\times$ are such that no quotient λ_i/λ_j (for $i \neq j$) is a root of unity, then there exists a place v of K such that, after renumbering, either*

$$|\lambda_1|_v \geq |\lambda_2|_v \geq |\lambda_3|_v \geq |\lambda_4|_v \quad \textit{with} \quad |\lambda_1|_v > |\lambda_3|_v,$$

or

$$|\lambda_1|_v = |\lambda_2|_v = |\lambda_3|_v > |\lambda_4|_v \quad \textit{with } v \textit{ Archimedean.}$$

Sketch. Suppose, for a contradiction, that no such place exists. Then at every Archimedean place all four absolute values $|\lambda_i|_v$ coincide, while at every non-Archimedean place three of them coincide and the fourth is no larger. Regard K as a subfield of $\mathbb{C}$, enlarged so as to be closed under complex conjugation. The Archimedean condition at the distinguished embedding then reads

$$\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2} = \lambda_3\overline{\lambda_3} = \lambda_4\overline{\lambda_4}.$$

Now let v be any non-Archimedean place, and let $\bar{v}$ denote its conjugate place, defined by $|x|_{\bar{v}} := |\overline{x}|_v$. The displayed equation gives $|\lambda_i|_v\,|\lambda_i|_{\bar{v}} = |\lambda_j|_v\,|\lambda_j|_{\bar{v}}$ for all i, j. Combined with the hypothesis that at v (and likewise at $\bar{v}$) three of the four absolute values coincide and the fourth is no larger, an elementary case analysis forces all four absolute values to coincide at v. Hence $|\lambda_i/\lambda_j|_v = 1$ for every place v of K, so that each quotient λ_i/λ_j has height zero and is therefore a root of unity, by Kronecker's theorem (Section 2.1)—contradicting the hypothesis. □

Combining Lemma 24 with the MSTV Theorem yields the current unconditional state of the art.

Theorem 25 (Mignotte–Shorey–Tijdeman, Vereshchagin, Bacik)**.** *The Skolem Problem is decidable for LRS of order at most 4 over the algebraic numbers; moreover, all zeros of such LRS can be effectively computed.*

Order 5 and beyond. At order 5, decidability of the Skolem Problem remains open, and the analogue of Lemma 24 for five numbers fails: a non-degenerate LRS of order 5 may have four dominant characteristic roots at every Archimedean place and three or more at every non-Archimedean place, thereby escaping the MSTV class. For rational LRS, the outstanding

hard configuration at order 5 comprises simple LRS whose dominant roots form two complex-conjugate pairs of equal modulus, together with one further real root of strictly smaller modulus. (An order-5 LRS with repeated roots has at most four distinct roots, and hence lies in the MSTV class; a place-hunting argument in the style of Lemma 24 disposes of the configuration in which all five roots share the same modulus [76].) We shall encounter this configuration again in Section 7 (in the Turing reduction from Skolem at order 5 to Positivity for simple LRS at order 10) and, in a closely related order-6 guise, in the hard instance analyzed in Section 8.3. Nevertheless, recent conditional results, which we discuss in Section 6, do establish decidability for simple LRS of order 7 or less, but only assuming the Exponential Local-Global Principle [76, 95], an unproven yet widely believed classical number-theoretic conjecture. We note in passing that for *reversible* LRS—integer LRS whose bi-infinite completion takes only integer values—decidability at order 7 or less is known unconditionally [76]; the Positivity Problem for this class is studied in [68].

5.4 Complexity at Low Orders

The decidability results of the preceding subsection rest on effective thresholds N with the property that any zero of the given LRS $\boldsymbol{u}$ must satisfy $n < N$. Tracing through the underlying Baker-type estimates reveals that these thresholds are at most exponential in the bit-size of the description of $\boldsymbol{u}$, and hence have polynomial bit-size themselves: this holds for non-degenerate LRS of order at most 4, and more generally for non-degenerate LRS of any fixed order in the MSTV class [43, 13]. Deciding the Skolem Problem on these classes therefore reduces, in polynomial time, to the Bounded Skolem Problem (Problem 21), and the complexity question becomes how efficiently one can search the exponentially large range $\{0, 1, \ldots, N\}$ for a zero.

A first answer is implicit in earlier work on the Orbit Problem [43]: the Bounded Skolem Problem, for arbitrary LRS, lies in $\mathsf{NP}^{\mathsf{RP}}$. Indeed, one may guess the index $n \leq N$—a polynomial number of bits—and verify the guess by constructing, via iterated matrix squaring in the companion form of Section 3.2, a polynomial-size arithmetic circuit representing the integer u_n, whose vanishing is an instance of the problem EqSLP of Section 2.4, decidable in randomized polynomial time [113]. In particular, the Skolem Problem for LRS of order at most 4 lies in $\mathsf{NP}^{\mathsf{RP}}$. In a similar vein, S. Akshay *et al.* established an exponential zero bound for LRS of arbitrary order all of whose characteristic roots are roots of real algebraic numbers, placing the Skolem Problem for this class in $\mathsf{NP}^{\mathsf{RP}}$ as well, alongside a matching NP-hardness lower bound [3].

These upper bounds were recently substantially improved as follows [13].

Theorem 26. *For every fixed $d \in \mathbb{N}$, the Bounded Skolem Problem for LRS of order at most d lies in* coRP. *Consequently, the Skolem Problem for LRS of order at most* 4 *lies in* coRP.

The same reasoning places the Skolem Problem in coRP for LRS of any fixed order in the MSTV class, as well as for the class of S. Akshay *et al.* at any fixed order. Moreover, the corresponding function problem—computing *all* zeros up to N, output as a finite set together with finitely many arithmetic progressions—is solvable in polynomial time with an oracle for EqSLP, and hence in randomized polynomial time. In effect, Theorem 26 shaves a full exponential off the earlier guess-and-check procedures.

The algorithm. The proof of Theorem 26 follows Skolem's p-adic method, in the form presented in Section 5.2, and turns it into an efficient search procedure. Choose a prime $p \geq d + 2$ not dividing the coefficient a_0 of the recurrence (1); by the fact that the product of all primes less than x grows exponentially in x, such a prime exists of magnitude polynomial in the bit-size of $\boldsymbol{u}$. As in Section 5.2, the LRS splits into polynomially many subsequences $\langle u_{Ln+r} \rangle_{n=0}^{\infty}$, each interpolated by a p-adic analytic function $f_r : \mathbb{Z}_p \to \mathbb{Z}_p$. A subsequence that is identically zero—a case detected by inspecting d initial terms—contributes a full arithmetic progression of zeros.

For the remaining subsequences, two further ingredients enter. First, a quantitative refinement of Skolem's argument, due to A. J. van der Poorten and H. P. Schlickewei [124], shows that each such f_r has at most $d-1$ zeros on the closed unit disc (where zeros are counted in the algebraically closed field $\mathbb{C}_p$ of Section 2.3). Second, by the p-adic Weierstrass Preparation Theorem (Section 2.3), whether a prescribed residue class modulo p^j contains such a zero is governed by the valuations of finitely many coefficients of a suitably recentered Mahler-series expansion of f_r; these coefficients are in turn computable from d terms of the LRS taken in arithmetic progression, evaluated modulo a prime power of polynomial bit-size via iterated matrix squaring, so that the search requires no p-adic arithmetic at all. The algorithm now performs a depth-first search of the residue classes modulo p^j, for successively larger j, retaining those classes that contain a zero of some f_r. At depth $j = O(\log N)$, each surviving class corresponds to at most one potential zero of $\boldsymbol{u}$ in $\{0, 1, \ldots, N\}$, and by the zero-count bound only polynomially many classes survive; the resulting indices are the *candidate* zeros of $\boldsymbol{u}$. Finally, $\boldsymbol{u}$ has a zero among the candidates if and only if the product of the corresponding terms vanishes—once again a polynomial-size arithmetic circuit—which is decided by a single invocation of the randomized EqSLP test. Randomness therefore enters only in this final step; the search itself is entirely deterministic.

In Section 8.3 we shall see these p-adic techniques in action on a concrete order-6 sequence lying outside the MSTV class: for that sequence the (unbounded) Skolem Problem remains open, yet zero-freeness throughout $\{0, 1, \ldots, 10^{1000}\}$ can be certified in practice in well under one minute.

We close with two remarks. First, the running time of the algorithm of Theorem 26 depends exponentially on the order d, as is to be expected: the NP-hardness reduction of Section 5.1 applies equally to the Bounded Skolem Problem, though it requires LRS of unbounded order. In fact, this exponential dependence only manifests itself through the number of subsequences L, which is itself bounded above by p^{d^2}.

Note also that the theorem imposes no non-degeneracy hypothesis: degenerate LRS are accommodated because the splitting into subsequences isolates the identically zero ones. Second, since the sole source of randomness is arithmetic-circuit identity testing, whose derandomization is a longstanding open problem (Section 2.4), an intriguing question is whether the coRP upper bounds above can be improved to deterministic polynomial time. We return to these complexity questions, and to the role of the Bounded Skolem Problem as a stepping stone toward decidability of the full Skolem Problem, in Section 9.3.

6 The Skolem Problem: Conditional and Alternative Approaches

We now turn to developments of the last decade, which extend the reach of the unconditional results of Section 5 in two complementary directions. The first is conditional: assuming one or more classical number-theoretic conjectures—most notably the Exponential Local-Global Principle and the p-adic Schanuel Conjecture, and, in more recent work, Cramér-type heuristics on prime gaps—the Skolem Problem becomes decidable, in particular for the entire class of simple LRS. The second direction alters the problem itself: one may seek p-adic rather than integer zeros (the p-adic Skolem Problem), restrict the set of indices at which zeros are sought (Universal Skolem Sets), or generalize the notion of zero (twisted rational zeros). The Bi-Skolem Problem (Problem 6), introduced in Section 1, plays a central role in the first two subsections.

6.1 The Bi-Skolem Problem and the Exponential Local-Global Principle

Recall the Bi-Skolem Problem (Problem 6): given an LRBS $\boldsymbol{u} = \langle u_n \rangle_{n=-\infty}^{\infty}$, decide whether there exists $n \in \mathbb{Z}$ with $u_n = 0$. The Bi-Skolem Problem reduces to the Skolem Problem: $\boldsymbol{u}$ has a zero if and only if at least one of the one-way infinite sequences $\langle u_n \rangle_{n=0}^{\infty}$ and $\langle u_{-n} \rangle_{n=0}^{\infty}$, both

of which are LRS, has a zero. Whether there is a reduction in the converse direction is open. Nevertheless, as we shall see, the Bi-Skolem Problem for simple LRBS is decidable assuming the conjecture that we now introduce.

Let $\boldsymbol{u}$ be a rational LRBS. Observe that $\boldsymbol{u}$ takes values in the ring $\mathbb{Z}[1/b]$ for a suitable non-zero integer b: indeed, it suffices that $\mathbb{Z}[1/b]$ contain the recurrence coefficients $a_0, \ldots, a_{d-1}$, the reciprocal a_0^{-1}, and the initial values $u_0, \ldots, u_{d-1}$, since then running the recurrence (1) forward and backward stays within $\mathbb{Z}[1/b]$. For any integer $m \geq 2$ with $\gcd(b, m) = 1$, the LRBS may therefore be reduced modulo m, and the resulting bi-infinite sequence in $\mathbb{Z}/m\mathbb{Z}$ is periodic, the companion matrix being invertible modulo m. In particular, the presence or absence of a zero in the reduction modulo m is decided by inspecting a single period; and if the reduction is everywhere non-zero, then certainly $u_n \neq 0$ for all $n \in \mathbb{Z}$. The following conjecture asserts that, for simple LRBS, this sufficient condition for zero-freeness is also necessary.

Conjecture 27 (Exponential Local-Global Principle)**.** *Let $\boldsymbol{u}$ be a simple rational LRBS taking values in $\mathbb{Z}[1/b]$. Then $\boldsymbol{u}$ is everywhere non-zero on $\mathbb{Z}$ if and only if there exists an integer $m \geq 2$ with $\gcd(b, m) = 1$ such that $\boldsymbol{u}$ is everywhere non-zero modulo m.*

The conjecture was first formulated by T. Skolem in 1937 [116], in the context of purely exponential Diophantine equations, and was historically known as the *Skolem Conjecture*; following [20], we adopt the more descriptive name. (Skolem's original formulation concerns such equations in several integer unknowns; Conjecture 27 is its specialization to a single exponent, which is the case relevant to LRBS.) As noted in Section 1, the truth of the conjecture immediately yields a decision procedure for the Bi-Skolem Problem on simple LRBS, by parallel search for either a zero or a witnessing modulus. On the other hand, it is not known whether the truth of the conjecture implies decidability of the Skolem Problem itself; as we shall see in the next subsection, bridging this gap is where the weak p-adic Schanuel Conjecture enters.

The restriction to simple LRBS cannot be dropped, as the following counterexample from [76] shows.

Example 28. Consider the LRBS $\boldsymbol{u}$ defined by $u_{n+2} = 4u_{n+1} - 4u_n$ with $u_0 = 1$, $u_1 = 6$, so that $u_n = (2n+1)\,2^n$ with the repeated characteristic root $\lambda = 2$. Here $\boldsymbol{u}$ takes values in $\mathbb{Z}[1/2]$, so the admissible moduli are the odd integers $m \geq 3$. The sequence is everywhere non-zero on $\mathbb{Z}$, since the polynomial factor $2n+1$ has no integer root. However, for any odd m, the congruence $2n + 1 \equiv 0 \pmod{m}$ has infinitely many solutions $n \in \mathbb{Z}$, and consequently $u_n \equiv 0 \pmod{m}$ for infinitely many n. Thus no admissible modulus witnesses the zero-freeness of $\boldsymbol{u}$.[3]

What is known unconditionally. The Exponential Local-Global Principle has been established in the following cases:

- certain families of LRBS of order 3, by A. Schinzel [110, 111];
- all simple LRBS whose characteristic roots generate a multiplicative group of rank one—in particular, all simple LRBS of order 2—by B. Bartolomé, Y. Bilu, and F. Luca [20];
- 'almost all' instances, in a suitable density-theoretic sense, by C. Bertók and L. Hajdu [27]; an extension to LRBS over number fields appears in [28].

In full generality the principle remains wide open, and constitutes an important challenge in number theory.

[3]For certifying zero-freeness over the domain $\mathbb{N}$ (as opposed to $\mathbb{Z}$), any integer modulus $m \geq 2$ could in principle be considered, since u_n is always an integer for $n \in \mathbb{N}$. Nevertheless, no witnessing modulus exists: writing $m = o \cdot 2^\ell$, with o odd, we see that there are infinitely many $n \in \mathbb{N}$ such that $u_n \equiv 0 \pmod{m}$, namely all those $n \geq \ell$ such that $2n + 1$ is divisible by o.

6.2 The p-adic Schanuel Conjecture and Conditional Decidability for Simple LRS

Granted the Exponential Local-Global Principle, the parallel search described in the preceding subsection already yields a conditional decidability result, which we record for later use.

Proposition 29. *Assume the Exponential Local-Global Principle (Conjecture 27). Then the Bi-Skolem Problem is decidable for simple LRBS.*

To progress from the Bi-Skolem Problem to the Skolem Problem proper, what is needed is the ability to compute *all* the zeros of an LRBS, rather than merely to decide whether one exists. This is supplied by a second classical conjecture, as shown in [31].

The p-adic Schanuel Conjecture. Schanuel's classical conjecture, a unifying principle in transcendence theory, asserts that for any complex numbers $\alpha_1, \ldots, \alpha_s$ that are linearly independent over $\mathbb{Q}$, the field $\mathbb{Q}(\alpha_1, \ldots, \alpha_s, e^{\alpha_1}, \ldots, e^{\alpha_s})$ has transcendence degree at least s over $\mathbb{Q}$. We shall require the following p-adic analogue.

Conjecture 30 (p-adic Schanuel Conjecture)**.** *Let $s \geq 1$, and let $t_1, \ldots, t_s \in \mathbb{C}_p$ be linearly independent over $\mathbb{Q}$, with $v_p(t_i) > 1/(p-1)$ for $1 \leq i \leq s$. Then*

$$\operatorname{trdeg}_{\mathbb{Q}} \mathbb{Q}\big(t_1, \ldots, t_s, \exp_p(t_1), \ldots, \exp_p(t_s)\big) \;\geq\; s,$$

where $\operatorname{trdeg}_{\mathbb{Q}}$ *denotes transcendence degree over* $\mathbb{Q}$.

For the results of this subsection, only the following weaker formulation is required (see [39, Conjecture 3.10]); Conjecture 30 will be needed in its full strength in Section 6.3.

Conjecture 31 (Weak p-adic Schanuel Conjecture)**.** *Let $\alpha_1, \ldots, \alpha_s \in 1 + p\mathbb{Z}_p$ be non-zero algebraic numbers whose logarithms $\log_p \alpha_1, \ldots, \log_p \alpha_s$ are linearly independent over $\mathbb{Q}$. Then these logarithms are algebraically independent over the algebraic numbers; that is, for every non-zero polynomial $P \in \mathbb{Q}_p[X_1, \ldots, X_s]$ whose coefficients are algebraic over $\mathbb{Q}$, it holds that $P(\log_p \alpha_1, \ldots, \log_p \alpha_s) \neq 0$.*

In the case where P is linear, the conjecture—the non-vanishing of $\beta_0 + \beta_1 \log_p \alpha_1 + \cdots + \beta_s \log_p \alpha_s$ for coefficients β_i algebraic over $\mathbb{Q}$ and not all zero—is a theorem of A. Brumer [36], the p-adic analogue of Baker's Theorem on linear independence of logarithms. The full algebraic-independence statement, like its Archimedean cousin, appears to lie well beyond the reach of current transcendence methods.

Computing all zeros. The following result of [31] supplies the bridge from the Bi-Skolem Problem to the Skolem Problem.

Theorem 32. *Assume the weak p-adic Schanuel Conjecture (Conjecture 31). Then the set of zeros of a non-degenerate LRBS can be computed using an oracle for the Bi-Skolem Problem. In particular, the Skolem Problem is Turing-reducible to the Bi-Skolem Problem.*

The procedure behind Theorem 32 is recursive. Given a non-degenerate LRBS $\boldsymbol{u}$—which we may assume is not identically zero, this case being trivial to detect—first query the oracle: if $\boldsymbol{u}$ is zero-free, there is nothing further to do. Otherwise, search systematically until a zero is found (the search must succeed, given the oracle's answer), and re-index so that the zero lies at index 0. The heart of the matter is now to compute a positive integer M such that $u_{Mn} \neq 0$ for all $n \in \mathbb{Z} \setminus \{0\}$. Granted this, the subsequence $\langle u_{Mn} \rangle_{n=-\infty}^{\infty}$ contains no zero beyond the one already found, and the remaining subsequences $\langle u_{Mn+r} \rangle_{n=-\infty}^{\infty}$, for $0 < r < M$, are again LRBS, on which the procedure is called recursively. Since each recursive call uncovers a fresh zero of

$\boldsymbol{u}$, of which there are only finitely many by the Skolem–Mahler–Lech Theorem, the recursion terminates.

The integer M is obtained via Skolem's p-adic method (Section 5.2), and it is here that the weak p-adic Schanuel Conjecture enters. Choose a prime p, not dividing a_0, for which all characteristic roots of $\boldsymbol{u}$ lie in $\mathbb{Z}_p$ (infinitely many primes have this property, by the Chebotarev density theorem), and let L be a positive integer such that $\lambda^L \equiv 1 \pmod{p}$ for every characteristic root λ. Then $u_{Ln} = f(n)$ for a p-adic power series $f(X) = \sum_{j \geq 0} a_j X^j$ whose coefficients a_j are polynomial expressions, with coefficients algebraic over $\mathbb{Q}$, in the p-adic logarithms $\log_p \lambda^L$ of a maximal multiplicatively independent subset of the characteristic roots (such a subset, together with the accompanying multiplicative relations, is computable thanks to a theorem of D. Masser [88]). These logarithms are linearly independent over $\mathbb{Q}$, so the weak p-adic Schanuel Conjecture guarantees that each a_j vanishes only when the corresponding polynomial expression vanishes identically. One can therefore identify the first non-zero coefficient a_{j_0} and compute its valuation $\nu := v_p(a_{j_0})$; a short ultrametric argument then shows that $M := L\,p^{\nu+1}$ has the required property. We refer to [31] for the details.

Combining Proposition 29 with Theorem 32—and noting that all LRBS arising in the recursion are simple whenever the input is—yields the main conditional result of this section.

Theorem 33. *Assume the Exponential Local-Global Principle (Conjecture 27) and the weak p-adic Schanuel Conjecture (Conjecture 31). Then the Skolem Problem is decidable for all simple LRS.*

At low orders, the weak p-adic Schanuel Conjecture can moreover be dispensed with: decidability of the Skolem Problem for simple LRS of order at most 5 holds assuming the Exponential Local-Global Principle alone [76], and likewise at orders 6 and 7 [95].

Two features of the procedure underlying Theorem 33 deserve emphasis. First, the two conjectures are needed only for *termination*: whenever the procedure halts, its output is unconditionally correct. Second, upon halting, the procedure produces a stand-alone certificate of that output: a partition of the input LRBS into finitely many subsequences, each containing at most one zero, in which every zero-free subsequence is accompanied by a witnessing modulus—a modular invariant in the sense of Section 4.3—and every subsequence containing a zero is accompanied by p-adic data certifying that this zero is unique. The certificate can be checked independently, without appeal to either conjecture.

The SKOLEM tool. The algorithm underlying Theorem 33—known as *leapfrogging*, after the ever-larger jumps M by which zeros are successively isolated—is implemented in the SKOLEM tool [31], available online at `https://skolem.mpi-sws.org/`. Given a simple integer LRS, the tool computes all zeros of its bi-infinite completion and produces the independent certificates described above; more recent versions incorporate further algorithms, including a Baker–Davenport-style procedure and the p-adic algorithm that we discuss in Section 6.3. In practice, the tool handles many interesting instances within seconds—including several examples from the literature to which no previously known technique applied—although random instances of order 7 and beyond typically remain out of practical reach [31].

6.3 The p-adic Skolem Problem

Recall from the proof of the Skolem–Mahler–Lech Theorem in Section 5.2 that, given a non-degenerate integer LRS $\boldsymbol{u}$ and a prime p not dividing the coefficient a_0 of its recurrence, there exist a positive integer L and analytic functions $f_r : \mathbb{Z}_p \to \mathbb{Z}_p$, for $0 \leq r < L$, such that $f_r(n) = u_{Ln+r}$ for all $n \in \mathbb{N}$. A *p-adic zero* of $\boldsymbol{u}$ is a zero $x \in \mathbb{Z}_p$ of one of the interpolants f_r. Every natural-number zero of $\boldsymbol{u}$ gives rise to a p-adic zero, and Strassmann's Theorem shows that a non-degenerate LRS has only finitely many p-adic zeros: this is the ineffective

finiteness argument at the heart of the Skolem–Mahler–Lech Theorem. Concretely, and without any p-adic terminology, $\boldsymbol{u}$ has a p-adic zero if and only if there exist natural numbers $n_1, n_2, \ldots$ with $u_{n_j} \equiv 0 \pmod{p^j}$ and $n_j \equiv n_{j+1} \pmod{p^j}$ for all j—a coherent family of zeros modulo ever-higher powers of p—and the corresponding p-adic zero arises from a natural-number zero precisely if the family can be taken eventually constant.

An LRS may well possess p-adic zeros that do not come from natural-number zeros, as the following example from [12] illustrates.

Example 34. Let $u_n = 4^n + 2$ and $p = 3$, and set $n_j := (1 + 3^j)/2$. An easy induction shows that $u_{n_j} \equiv 0 \pmod{3^{j+1}}$ for all j, and $n_j \to \frac{1}{2}$ in $\mathbb{Z}_3$; consequently the interpolant f (here $L = 1$) satisfies $f(\frac{1}{2}) = 0$. Thus $\boldsymbol{u}$, which clearly has no natural-number zero, has the 3-adic zero $\frac{1}{2}$. This zero is moreover *rational*, reflecting the identity $4^{1/2} + 2 = 0$ for the choice $4^{1/2} := -2$ of square root.

In general, the p-adic zeros of an LRS may be natural numbers, rational numbers (in the sense of the example), or irrational; and, as we discuss below, no effective means is known of telling these cases apart. The following problem is nonetheless of both theoretical and practical significance.

Problem 35 (The p-adic Skolem Problem)**.** *Given a non-degenerate LRS $\boldsymbol{u}$ and a prime p not dividing the coefficient a_0 of its recurrence, decide whether $\boldsymbol{u}$ has a p-adic zero.*

The corresponding function problem asks to compute exact representations of *all* p-adic zeros of $\boldsymbol{u}$: each zero is output as a disc in $\mathbb{Z}_p$ containing that zero and no other, from which it may then be approximated to any prescribed precision—much as an algebraic number is conventionally represented by its minimal polynomial together with an isolating approximation (Section 2.1). The following is the main result of [12].

Theorem 36. *Assume the p-adic Schanuel Conjecture (Conjecture 30). Then the p-adic Skolem Problem is decidable, and exact representations of all p-adic zeros of a given LRS can be computed. As in Section 6.2, the conjecture is needed only for termination: the output of the algorithms is unconditionally correct.*

The algorithm searches the residue classes modulo p^j, for successively larger j, in the spirit of Section 5.4. A residue class on which an interpolant is non-vanishing modulo p^j is discarded; a class on which the hypothesis of Hensel's Lemma is satisfied is certified to contain exactly one zero, which Newton iteration then approximates to any desired precision; the remaining classes are refined further. Throughout, a Newton-polygon count of the zeros in each class (a generalization of Strassmann's Theorem) certifies when all zeros have been accounted for. If every zero encountered is simple, this procedure terminates; the obstruction is zeros of higher multiplicity, on which the Hensel test can never succeed. To deal with these, the interpolants are first factored into irreducible exponential polynomials, and the search is conducted on each irreducible factor separately. The key technical result of [12] is that, assuming the p-adic Schanuel Conjecture and under suitable hypotheses, all common p-adic zeros of two *coprime* exponential polynomials are rational; applied to an irreducible factor and its derivative, this shows that the multiple zeros of each factor are rational. The algorithm therefore enumerates rational numbers in parallel with the residue search—guided, in practice, by the p-adic approximations already computed—and thereby locates every zero. (An in-principle decision procedure for the p-adic Skolem Problem can also be extracted from Mariaule's conditional decidability of the first-order theory of $(\mathbb{Z}_p; +, \cdot, \exp_p)$ [86]; the algorithm just described is, by contrast, designed to be implementable.)

Why does all this not resolve the Skolem Problem itself? Given the exact representation of a p-adic zero, one can approximate it to arbitrary precision, and thereby certify that it *is* a particular natural number when this happens to be the case; but no effective procedure is known

that can certify a p-adic zero to be irrational, or even to be a non-integer. Whether the Skolem Problem reduces to the p-adic Skolem Problem is thus itself an intriguing open question.

The Simultaneous Skolem Problem. The technical result on coprime exponential polynomials has a notable consequence for the following natural variant of the Skolem Problem.

Problem 37 (The Simultaneous Skolem Problem)**.** *Given two LRS $\boldsymbol{u}$ and $\boldsymbol{v}$, does there exist $n \in \mathbb{N}$ such that $u_n = v_n = 0$?*

Theorem 38. *Assume the p-adic Schanuel Conjecture (Conjecture 30). Then the Simultaneous Skolem Problem is decidable for* coprime *LRS, that is, for pairs of LRS whose exponential-polynomial representations share no non-trivial common factor [12].*

Indeed, a common natural-number zero of $\boldsymbol{u}$ and $\boldsymbol{v}$ is in particular a common p-adic zero of the associated coprime exponential polynomials, and by the technical result above all such common zeros are rational. Since the algorithm of Theorem 36 computes exact representations of the finitely many p-adic zeros of each sequence, the common zeros can be identified exactly—the accompanying rational search is guaranteed to reach each of them—and then tested for membership in $\mathbb{N}$.

Unconditional results on the Simultaneous Skolem Problem, in the variant where several LRS are considered at once, were recently obtained in [14]: given k linearly independent non-degenerate algebraic LRS satisfying a common recurrence of order d, the set of common zeros is computable whenever $d - k \leq 2 \log_3 d$. The essential observation is that a common zero of several LRS is a zero of *every* linear combination of them, and the combination may be chosen so as to make the resulting LRS as tractable as possible. The same paper establishes a tight correspondence between the Simultaneous Skolem Problem and the Subspace Orbit Problem, to which we return in Section 9.2.

The function-problem algorithm of Theorem 36 is implemented in the SKOLEM tool, in the case where all characteristic roots lie in $\mathbb{Z}_p$ (which, as noted in Section 6.2, holds for infinitely many primes). It is this p-adic mode that underlies the practical certifications of zero-freeness over very large ranges mentioned in Section 5.4; we shall see it in action on a concrete order-6 instance in Section 8.3, whose sixteen 17-adic zeros are precisely coherent families of zeros modulo all powers of 17 in the sense described above.

6.4 Universal Skolem Sets

A further approach to the Skolem Problem, initiated in [79], restricts the domain in which zeros are sought instead of placing structural restrictions on the LRS under consideration.

Definition 39. An infinite, recursive set $\mathcal{S} \subseteq \mathbb{N}$ is a *Universal Skolem Set* if there is a procedure that, given an LRS $\boldsymbol{u}$, decides whether there exists $n \in \mathcal{S}$ with $u_n = 0$.

Decidability of the Skolem Problem is thus equivalent to the assertion that $\mathbb{N}$ itself is a Universal Skolem Set, and one may hope to make progress by exhibiting Universal Skolem Sets of ever-higher density.

The known constructions share, broadly speaking, a common mechanism, which one might call a *transfer principle*: for indices n lying in the set $\mathcal{S}$ under construction, a hypothetical zero $u_n = 0$ forces a large number of auxiliary sequences—obtained from $\boldsymbol{u}$ by suitable arithmetic twists—to vanish at small indices. Effective bounds on the *number* of zeros of such sequences, in the vein of Theorem 17, are thereby brought to bear on the *magnitude* of n: beyond an effectively computable threshold, a zero of $\boldsymbol{u}$ in $\mathcal{S}$ would violate the zero-count bounds. The Skolem Problem on $\mathcal{S}$ consequently reduces to checking finitely many indices. We shall see this mechanism concretely in the next subsection.

The known constructions. The original paper [79] exhibited a Universal Skolem Set defined by a sparse super-exponential recursion: with $f(n) := \lfloor \sqrt{\log n} \rfloor$, set $s_0 := 1$ and inductively $s_n := n! + s_{f(n)}$ for $n > 0$; then $\mathcal{S}_0 := \{s_n : n \in \mathbb{N}\}$ is a Universal Skolem Set. The construction has the advantage of being elementary, but the resulting set has density zero.

A subsequent construction [81] achieved positive lower density, at the price of restricting universality to simple LRS: the deciding procedure is guaranteed to exist for simple sequences only. Here the requisite zero-count bounds are drawn from the theory of exponential Diophantine equations, in the tradition of Evertse's bounds on the number of solutions of S-unit equations, combined with sieve methods.

Both limitations were overcome in [78], which constructs a single Universal Skolem Set $\mathcal{S}_1$, valid for arbitrary LRS, whose lower density exceeds 0.29 unconditionally; moreover, subject to the Bateman–Horn Conjecture in number theory, the very same set $\mathcal{S}_1$ has density one. The Bateman–Horn Conjecture predicts the asymptotic frequency at which a finite family of irreducible integer polynomials (satisfying the evident necessary conditions) simultaneously assumes prime values; it subsumes, for example, a quantitative form of the twin prime conjecture. Its role in [78] is to furnish an abundance of primes of the particular shape required by the transfer mechanism.

A Universal Skolem Set of density one does in fact exist *unconditionally*: this emerges from the analysis of *large zeros* of LRS presented in the next subsection. The unconditional construction is, however, of a rather different nature: the set in question is defined by direct reference to the zeros of LRS themselves, whereas $\mathcal{S}_1$ is defined by intrinsic arithmetic conditions, making the latter construction both more demanding and arguably more informative. In particular, the unconditional result complements, rather than supersedes, the Bateman–Horn-conditional statement above: it does not remove the need for the conjecture in establishing that $\mathcal{S}_1$ itself has density one. We note finally that Universal Skolem Sets are closed under finite unions and under shifts; a diverse supply of such sets is therefore of independent value, since one might hope to establish that $\mathbb{N}$ is a Universal Skolem Set by expressing it as a finite union of (possibly shifted) Universal Skolem Sets.

6.5 Large Zeros and the Cramér Heuristic

The developments of this subsection, drawn from [83] (which improves upon the earlier paper [82]), revolve around the following question: *how large can the zeros of an LRS be*, measured against the size of the data defining the sequence? For an integer LRS $\boldsymbol{u}$ of order d, define its *height* $H_{\boldsymbol{u}}$ to be the maximum of d and the absolute values of the recurrence coefficients and initial terms (note that this measures the *magnitude* of the data, rather than its bit-size). Call a zero $u_n = 0$ of a non-degenerate LRS of sufficiently large height a *large zero* if

$$n \;\geq\; \exp\exp\bigl(10\, H_{\boldsymbol{u}} \log H_{\boldsymbol{u}}\bigr). \tag{10}$$

Plainly, if one could show that large zeros do not exist, then every zero would admit a doubly exponential bound in the height, and decidability of the Skolem Problem would follow by direct search. No such proof is currently in sight; but there are good heuristic grounds—and, as we shall see, unconditional density results—supporting the view that large zeros should not exist. Indeed, no family of non-degenerate LRS is known whose largest zeros grow even singly exponentially in the height.

Good primes. The engine of the analysis is a Frobenius-twisting argument, in the spirit of the transfer principle of the preceding subsection. Suppose $u_n = 0$, and write $n = P + m$ where P is a prime not much smaller than n and m is (comparatively) small. Reducing the exponential-polynomial expression for u_n modulo a prime ideal above P, the Frobenius automorphism transforms the relation $u_n = 0$ into the assertion that P divides the norm of a

twisted version of the expression, evaluated at the small index m: the characteristic roots are permuted, and the exponent n is replaced by m. Call the prime P *bad* if it divides such a norm for a non-zero twisted expression; in [83] it is shown, by counting considerations, that the bad primes form a very sparse set (of counting function $O(X^{2/3})$), so that the remaining *good* primes have density one among all primes. For a good prime P, the twisted expression must actually vanish at m; and the number of zeros of each twisted sequence is controlled, via the Amoroso–Viada bounds underlying Theorem 17, by a quantity doubly exponential in $H_{\boldsymbol{u}}$. Since distinct primes P with $n = P + m$ yield distinct indices m, the upshot is that an interval of length roughly $n^{1/4}$ below a large zero n can contain at most doubly exponentially many (in $H_{\boldsymbol{u}}$) good primes—an extreme scarcity, given that good primes have density one among all primes.

Conditional decidability. The classical Cramér–Granville Conjecture asserts that gaps between consecutive primes satisfy $p_{j+1} - p_j = O\big((\log p_j)^2\big)$; it is motivated by the probabilistic model, going back to H. Cramér, in which the primes behave like a random set of integers of density prescribed by the prime number theorem, and it is supported by extensive computational evidence (though the best unconditional upper bound on prime gaps, $O(p_j^{0.525})$, is very far from it). Since good primes have density one among all primes, it is natural to posit the same gap bound for consecutive *good* primes; this is the strengthened form of the conjecture introduced in [83]. Granted it, every interval of length $n^{1/4}$ below n contains on the order of $n^{1/4}/(\log n)^2$ good primes—overwhelmingly more than the doubly exponential ceiling just described—and large zeros therefore cannot exist.

Theorem 40. *Assume the strengthened Cramér–Granville Conjecture of [83]. Then large zeros do not exist: for every non-degenerate LRS $\boldsymbol{u}$ of sufficiently large height, every zero satisfies $n < \exp\exp(10\, H_{\boldsymbol{u}} \log H_{\boldsymbol{u}})$. Consequently, the Skolem Problem is decidable.*

We note that the decidability conclusion is non-constructive in one mild respect: there are only finitely many LRS of height below the (unknown) absolute constant implicit in "sufficiently large", so *some* correct decision procedure for these instances exists, although the argument does not exhibit one. In any event, the doubly exponential magnitudes involved place the resulting algorithm well beyond practical use; the significance of Theorem 40 is conceptual, tying the Skolem Problem to a well-established heuristic model of the primes.

Unconditional consequences. Remarkably, the same machinery yields unconditional results. Consider the set

$$\mathcal{L} \;:=\; \big\{ n \in \mathbb{N} \;:\; \text{some non-degenerate LRS } \boldsymbol{u} \text{ has a large zero at } n \big\},$$

the set of natural numbers that arise as large zeros at all. Using a result of C. Jia on primes in short intervals in place of the conjectured gap bounds, the following is established in [83].

Theorem 41. *The set $\mathcal{L}$ has density zero; in fact, its counting function satisfies $\#\big(\mathcal{L} \cap [1, X]\big) = O\big(X/(\log X)^B\big)$ for every constant $B > 0$.*

Corollary 42. *The set $\mathcal{S} := \mathbb{N} \setminus \mathcal{L}$ is a Universal Skolem Set of density one.*

Indeed, $\mathcal{S}$ is recursive; it has density one by Theorem 41; and it is universal because, by construction, the zeros in $\mathcal{S}$ of any non-degenerate LRS are bounded by the doubly exponential threshold (10), and may therefore be found by direct search. What remains open, of course, is whether one can take $\mathcal{S} = \mathbb{N}$ itself—that is, the Skolem Problem.

6.6 Twisted Rational Zeros

Our final topic on the Skolem side generalizes the very notion of a zero and, in doing so, explains the rational p-adic zeros encountered in Section 6.3. The starting point is a phenomenon concerning p-adic valuations of LRS, first observed by D. Marques and T. Lengyel for the Tribonacci sequence [87], defined by the recurrence $T_{n+3} = T_{n+2} + T_{n+1} + T_n$ with $T_0 = 0$ and $T_1 = T_2 = 1$, and regarded here as an LRBS: on suitable residue classes, the 2-adic valuation $v_2(T_n)$ is given by an explicit formula, either constant or of the shape $v_2(n-a)+c$ for a constant c, with $a \in \{0, -1, -4, -17\}$—and these four numbers are precisely the zeros of the bi-infinite Tribonacci sequence. This is no accident: every zero a of an LRBS gives rise, for all but finitely many primes p, to a formula of this shape, valid on a suitable residue class [33].

Such valuation formulas were studied systematically, for arbitrary primes p, in [32], whose main result is an algorithm that determines, for any given p, whether formulas of this kind exist for the Tribonacci sequence, and computes them when they do. More surprising is that such p-adic valuation formulas may also be centered at points where the sequence does *not* vanish. For instance, the LRBS $u_n = 2^n + 1$ has no zero whatsoever, yet for every prime $p \equiv \pm 3 \pmod 8$ the valuation $v_p(u_n)$ obeys a formula of the above shape centered at $a = 0$, on a suitable residue class. The explanation is that 0 is a 'hidden' zero of the sequence: one has $1 \cdot 2^0 + (-1) \cdot 1^0 = 0$, that is, the exponential-polynomial expression for u_n vanishes at $n = 0$ once its terms are *twisted* by suitable roots of unity.

These observations are captured by the following definitions from [33]. Let $\boldsymbol{u}$ be an LRBS with exponential-polynomial representation $u_n = \sum_{j=1}^{m} Q_j(n)\, \lambda_j^n$. A rational number $a = c/b$ is a *rational zero* of $\boldsymbol{u}$ if

$$Q_1(a)\, \lambda_1^a + \cdots + Q_m(a)\, \lambda_m^a \;=\; 0$$

for *some* choice of the b-th roots $\lambda_j^a := \sqrt[b]{\lambda_j^c}$, and a *twisted rational zero* if, moreover, each term may first be multiplied by an arbitrary root of unity ξ_j. Integer zeros are twisted rational zeros with trivial twists, and we have already met a genuinely rational zero in Example 34: there $a = \frac{1}{2}$ and $u_n = 4^n + 2$, with the requisite choice of square root being $4^{1/2} := -2$—a twist of $\sqrt{4} = 2$ by the root of unity -1.

The connection with p-adic zeros runs as follows: if, for some prime p outside a finite set determined by the recurrence, the p-adic valuation $v_p(u_{n_k})$ grows unboundedly along a sequence of integers n_k converging p-adically to a rational number a, then a must be a twisted rational zero of $\boldsymbol{u}$ [33]. (In the language of Section 6.3, this hypothesis is what a coherent family of zeros modulo increasing powers of p, converging to a rational point, provides.) Twisted rational zeros thus account for the rational points at which an LRS can become p-adically small without vanishing, and identifying a rational p-adic zero as a twisted rational zero separates the proper zeros of the sequence from the spurious ones.

The principal result of [33] extends the Skolem–Mahler–Lech Theorem to this generalized notion of zero.

Theorem 43. *A non-degenerate LRS over a field of characteristic* 0 *admits at most finitely many twisted rational zeros.*

The proof is a variation of an argument of M. Laurent: one first bounds the denominators of all twisted rational zeros—effectively so, when the LRS is defined over a number field—after which the statement reduces to the classical Skolem–Mahler–Lech Theorem together with known results on equations in roots of unity. Exactly as with the Skolem–Mahler–Lech Theorem itself, the finiteness is ineffective in the numerators: computing the full set of twisted rational zeros is not known to be possible in general, although, as with Skolem–Mahler–Lech, it can be carried out in many special cases. For the Tribonacci sequence, the twisted rational zeros are shown in [33] to be exactly

$$0, \quad -1, \quad -4, \quad -17, \quad \tfrac{1}{3}, \quad -\tfrac{5}{3}.$$

The first four are the integer zeros of the bi-infinite Tribonacci sequence, while the two genuinely rational values manifest themselves p-adically: for infinitely many primes p, one has $v_p(T_n) \geq v_p(n - \frac{1}{3})$ on a suitable residue class, and likewise for $-\frac{5}{3}$ [32].

7 The Positivity and Ultimate Positivity Problems

We now turn to the Positivity and Ultimate Positivity Problems. The flavor of this part of the survey differs markedly from that of Sections 5 and 6: the modular and p-adic certificates that drive much of the contemporary work on the Skolem Problem are of no use for Positivity—neither modular reduction nor p-adic absolute values are compatible with the ordering of the reals—and no analogue of the Universal-Skolem-Set program has been developed. What remains, in the vocabulary of Section 4, are Archimedean separation bounds, which must now control signs rather than mere magnitudes, and semialgebraic invariants. Decidability is accordingly known only at low orders and for special classes; moreover, as we shall see, at order 6 both the Positivity and the Ultimate Positivity Problems run into concrete obstructions from Diophantine approximation.

7.1 Statements and Basic Facts

We restate the two problems from Section 1 for convenience; recall that, throughout, 'positivity' means non-negativity.

Problem 44 (The Positivity Problem)**.** *Given an LRS $\boldsymbol{u}$ over $\mathbb{Q}$, decide whether $u_n \geq 0$ for all $n \in \mathbb{N}$.*

Problem 45 (The Ultimate Positivity Problem)**.** *Given an LRS $\boldsymbol{u}$ over $\mathbb{Q}$, decide whether there exists $N \in \mathbb{N}$ such that $u_n \geq 0$ for all $n \geq N$.*

A classical framing, originating in the formal-power-series literature [108, 107], groups these together with the problems of Section 5 into four basic decision problems for a given LRS: does $u_n = 0$ hold for some n (the Skolem Problem)? for infinitely many n? does $u_n \geq 0$ hold for all n (Positivity)? for all but finitely many n (Ultimate Positivity)? The second of these questions, in contrast to the other three, has long been known to be decidable [25].

Theorem 46 (Berstel–Mignotte)**.** *There is an effective procedure that, given an LRS $\boldsymbol{u}$ over $\mathbb{Q}$, decides whether $\boldsymbol{u}$ has infinitely many zeros.*

Indeed, by the Skolem–Mahler–Lech Theorem, $\boldsymbol{u}$ has infinitely many zeros if and only if some subsequence in its effective non-degenerate decomposition (Proposition 14) is identically zero, a condition that is checked by inspecting finitely many initial terms.

As shown in Section 1, decidability of Positivity entails decidability of the Skolem Problem, via the Hadamard-square construction, with at most a quadratic increase in order (but no such reduction is known to the Ultimate Positivity Problem). A more careful accounting of orders in reductions of this kind underlies the Turing reduction, presented in Section 7.4, of the Skolem Problem at order 5 to Positivity for simple LRS of order 10. We also recall from Section 1 that the positivity set $\{n \in \mathbb{N} : u_n \geq 0\}$ of any LRS has a well-defined natural density, which is computable to arbitrary precision [22, 67]; but this offers no purchase on the Positivity Problems themselves, which turn on sign violations that may be far too sparse for any density to detect.

A further basic fact, established in [63], is that the Positivity Problem over the real algebraic numbers is no harder than over the integers.

Theorem 47. *The Positivity Problem for LRS over $\mathbb{R} \cap \overline{\mathbb{Q}}$ Turing-reduces to the Positivity Problem for integer LRS. Moreover, the reduction preserves simplicity: for simple real algebraic LRS, an oracle for Positivity of simple integer LRS suffices.*

Unlike the norm construction of Proposition 20, which preserves zero sets outright, the reduction here is a genuine Turing reduction: the oracle is invoked on auxiliary integer LRS assembled from effective rational LRS approximations of the algebraic data, and the order of these auxiliary sequences is, in general, exponential in the order of the input and the degree of its field of definition.

As for lower bounds, both Positivity and Ultimate Positivity are coNP-hard: the LRS arising in the NP-hardness construction for the Skolem Problem (Section 5.1) are periodic, and on periodic sequences positivity and ultimate positivity coincide, so that the Hadamard-square reduction applies to both problems alike [99].

7.2 Decidability at Low Orders

The unconditional state of the art at low orders is the following [99], extending earlier results at orders up to 3: Ultimate Positivity was settled at order 2 by J. R. Burke and W. A. Webb [37], and for order-3 LRS with repeated roots by K. Nagasaka and J.-S. Shiue [93], while Positivity was settled at order 2 by V. Halava, T. Harju, and M. Hirvensalo [57] and at order 3 by V. Laohakosol and P. Tangsupphathawat [72].

Theorem 48. *The Positivity and Ultimate Positivity Problems are decidable for LRS of order at most* 5*. Moreover, Ultimate Positivity at order at most* 5 *is decidable in polynomial time, while Positivity at order at most* 5 *lies in* $\mathsf{coNP}^{\mathsf{PosSLP}}$*, a class contained in the Counting Hierarchy, which in turn sits within* PSPACE*.*

The analysis begins with a classical observation (see, e.g., [22]).

Proposition 49. *A non-zero LRS with no positive real dominant characteristic root has infinitely many negative terms (as well as infinitely many positive ones).*

Proposition 49 disposes at once of every configuration lacking a positive real dominant root: such an LRS is neither positive nor ultimately positive. In particular, the configuration of four dominant complex roots together with a real root of smaller modulus—the outstanding hard case for the *Skolem* Problem at order 5 (Section 5.3)—is trivial for Positivity. In the remaining configurations there is a positive real dominant root ρ; since non-degeneracy rules out $-\rho$ being a root as well, the dominant roots consist of ρ together with some number of complex-conjugate pairs, so that their count is odd: 1, 3, or 5 at order at most 5.

Normalizing, write

$$\frac{u_n}{\rho^n} \;=\; A(n) \;+\; \sum_{i=1}^{m}\Big(C_i(n)\,\lambda_i^n \;+\; \overline{C_i(n)}\,\overline{\lambda_i}^{\,n}\Big) \;+\; r(n), \tag{11}$$

where the λ_i are the quotients by ρ of the dominant complex roots—algebraic numbers of modulus 1 which, by non-degeneracy, are not roots of unity—and $r(n)$, the contribution of the non-dominant roots, decays exponentially. Repeated dominant roots make $A(n)$ or the $C_i(n)$ non-constant polynomials, but at order at most 5 such configurations either resolve immediately or reduce to the constant-coefficient analysis: a non-constant leading term in $A(n)$ forces the asymptotic sign, while a non-constant $C_i(n)$ produces oscillations of linearly growing amplitude that preclude positivity outright. The delicate cases thus have constant coefficients, and we illustrate two of them.

The Baker case: three dominant roots. Suppose the dominant roots are ρ and the complex pair $\rho\lambda_1, \rho\overline{\lambda_1}$. Writing $A(n) = c$, $C_1(n) = c_1 = |c_1|\,e^{i\varphi}$ and $\lambda_1 = e^{i\theta}$, Equation (11) becomes

$$\frac{u_n}{\rho^n} \;=\; c + 2\,|c_1|\cos(n\theta + \varphi) + r(n).$$

If $c > 2|c_1|$, then $\boldsymbol{u}$ is non-negative from an effectively computable index onward. If $c < 2|c_1|$ then, as θ is not a rational multiple of π, the sequence $\langle n\theta \bmod 2\pi\rangle$ equidistributes, the cosine term drops below $-c/(2|c_1|)$ infinitely often, and $\boldsymbol{u}$ is neither positive nor ultimately positive. In the critical case $c = 2|c_1|$, everything hinges on how fast $\cos(n\theta+\varphi)$ can approach -1, that is, on how well $n\theta+\varphi$ can approximate odd integer multiples of π. Baker's Theorem (in the sharp form of A. Baker and G. Wüstholz) supplies the answer: the distance in question is bounded below by an inverse polynomial in n, whence $c+2|c_1|\cos(n\theta+\varphi)$ is likewise bounded below by an inverse polynomial in n. Since $|r(n)|$ decays exponentially, we conclude that $u_n \geq 0$ beyond an effectively computable threshold N, and it remains to examine the initial segment $u_0, \ldots, u_N$.

The torus case: five dominant roots. At order 5 with five dominant roots there are no non-dominant roots at all, so that $r(n)$ vanishes identically and

$$\frac{u_n}{\rho^n} \;=\; c \;+\; c_1\lambda_1^n + \overline{c_1}\,\overline{\lambda_1}^{\,n} \;+\; c_2\lambda_2^n + \overline{c_2}\,\overline{\lambda_2}^{\,n} \;=\; c + h\big(\lambda_1^n, \lambda_2^n\big)$$

for an explicit real-valued function h of two unimodular arguments. The multiplicative relations holding between λ_1 and λ_2 can be computed, by a theorem of D. Masser [88]; they carve out a subgroup T of the torus of unimodular pairs, and Kronecker's theorem on simultaneous Diophantine approximation guarantees that the orbit $\langle(\lambda_1^n, \lambda_2^n)\rangle_{n=0}^{\infty}$ is *dense* in T. Consequently $\boldsymbol{u}$ is positive if and only if it is ultimately positive, if and only if $c + h \geq 0$ holds throughout T; and this last assertion can be rephrased as a sentence of the first-order theory of the reals in a fixed number of variables, decidable in polynomial time by Renegar's quantifier-elimination procedure [105].

Complexity. The remaining configurations are handled by variations on these themes, and the complexity claims of Theorem 48 can now be explained. For Ultimate Positivity, all the requisite computations—bases of multiplicative relations, Renegar sentences, Baker-type thresholds—can be carried out in polynomial time. For Positivity, one first decides Ultimate Positivity; in the affirmative case, the analysis moreover yields an effective threshold N, of magnitude at most exponential in the bit-size of $\boldsymbol{u}$, beyond which all terms are non-negative. Failure of Positivity is then equivalent to the existence of an index $n \leq N$ with $u_n < 0$: such an n has polynomially many bits and can be guessed, and the sign of u_n can be tested by encoding u_n—via the matrix form of Section 3.2 and iterated squaring—as an arithmetic circuit, that is, as an instance of the problem PosSLP of Section 2.4. This places the complement of Positivity in $\mathsf{NP}^{\mathsf{PosSLP}}$; since PosSLP lies in the Counting Hierarchy [8], so does Positivity, and the Counting Hierarchy is itself contained in polynomial space.

The order-6 wall. At order 6 a new configuration appears, in which $A(n)$ and $C_1(n)$ in (11) are *both* non-constant (linear) polynomials—as happens when the characteristic roots are a real number and a complex-conjugate pair, all of the same modulus and each of multiplicity two. Here the interplay between the two growing terms defeats the analysis above, and this is no shortcoming of the method: as we shall see in Section 7.4, precisely such sequences witness that deciding Positivity or Ultimate Positivity at order 6 would resolve longstanding open problems in Diophantine approximation.

7.3 The Simple Case

For simple LRS, the state of the art reaches considerably further [98, 100].

Theorem 50. *The Positivity Problem is decidable for simple LRS of order at most* 9*, with complexity in* $\mathsf{coNP}^{\mathsf{PosSLP}}$*, as in Theorem 48. The Ultimate Positivity Problem is decidable for simple LRS of arbitrary order, in polynomial space, and in polynomial time for each fixed order.*

The contrast between the two halves of the theorem is striking: Ultimate Positivity is settled unconditionally at *all* orders, whereas Positivity proper remains open beyond order 9. As we now explain, the two results rest on different number-theoretic inputs, and the divergence between them is no accident.

Ultimate Positivity at all orders. Let $\boldsymbol{u}$ be a simple non-degenerate LRS. By Proposition 49 we may assume a positive real dominant root ρ, and normalizing as in (11) we write

$$\frac{u_n}{\rho^n} \;=\; c \;+\; h\big(\lambda_1^n, \ldots, \lambda_k^n\big) \;+\; r(n),$$

where—simplicity being the crucial point—all coefficients are constants: h is an explicit real-valued linear function of k unimodular arguments and their conjugates, and $r(n)$ decays exponentially. As in the torus case of Section 7.2, one computes, via Masser's theorem [88], the multiplicative relations holding among $\lambda_1, \ldots, \lambda_k$ (equivalently, the integer linear relations among their arguments $\theta_1, \ldots, \theta_k$ and 2π); these determine a subtorus T in which the orbit $\langle(\lambda_1^n, \ldots, \lambda_k^n)\rangle_{n=0}^{\infty}$ is dense, by Kronecker's theorem. Renegar's procedure then determines which of three situations arises. If $c + h > 0$ throughout T then, T being compact, u_n/ρ^n is eventually bounded away from zero from below, and $\boldsymbol{u}$ is ultimately positive. If $c + h < 0$ somewhere on T then, by density, $u_n < 0$ infinitely often, and $\boldsymbol{u}$ is not ultimately positive. The delicate case is that in which $\min_T (c + h) = 0$. Then the dominant part

$$w_n \;:=\; \rho^n\big(c + h(\lambda_1^n, \ldots, \lambda_k^n)\big)$$

is a simple LRS taking only non-negative values, and the question is whether it eventually dominates the exponentially smaller tail $\rho^n\, r(n)$. Here a growth theorem of J.-H. Evertse and of A. J. van der Poorten and H. P. Schlickewei [53, 123] comes to the rescue: if $\boldsymbol{w}$ is non-zero then, for any ϱ smaller than the spectral radius ρ of $\boldsymbol{w}$, one has $|w_n| > \varrho^n$ for all sufficiently large n. Choosing ϱ strictly between the moduli of the non-dominant roots of $\boldsymbol{u}$ and ρ, the dominant part eventually outstrips the tail, and $\boldsymbol{u}$ is again ultimately positive. (If instead $\boldsymbol{w}$ is identically zero, then $\boldsymbol{u}$ coincides with its non-dominant part, and one recurses at lower order.)

Effective versus ineffective. The procedure just described is genuinely non-constructive. The growth theorem guarantees that w_n eventually dominates the tail, but its known proofs rest on the Subspace Theorem (in Schlickewei's p-adic form), which is ineffective: no bound is provided on the index from which the domination—and hence ultimate positivity—takes hold. This is not a removable blemish. For simple LRS of any fixed order, the ability to compute such threshold indices would immediately decide Positivity as well, since it would remain only to check the sign of an initial segment directly; and decidability of Positivity for simple LRS would, in view of the reductions presented in the next subsection, resolve the Skolem Problem at order 5. By contrast, the order-9 Positivity result of [98] runs on *effective* Diophantine bounds furnished by Baker's Theorem, deployed much as in Section 7.2; but the analytic and geometric arguments that bring Baker's Theorem to bear are confined to low orders, and do not extend beyond order 9. The two halves of Theorem 50 thus embody an instructive trade-off: the effective tools reach only low orders, while the tools that reach all orders are ineffective.

On the lower-bound side, both problems for simple LRS admit a polynomial-time reduction from $\forall\mathbb{R}$, the decision problem for the universal theory of the reals [100]. The latter is coNP-hard and contained in PSPACE, so that Ultimate Positivity for simple LRS is pinned between $\forall\mathbb{R}$ and PSPACE.

7.4 Diophantine Hardness

Why does decidability stall at order 6? The results of this subsection, established in [99], show that the obstruction is not a want of ingenuity: deciding Positivity or Ultimate Positivity at or-

der 6 would carry substantial consequences in the Diophantine approximation of transcendental numbers.

Two classical quantities measure how well a real number x can be approximated by rationals. The *Lagrange constant* $L_\infty(x)$ is the infimum of all $c > 0$ such that

$$\left|x - \frac{n}{m}\right| \;<\; \frac{c}{m^2}$$

admits infinitely many solutions $(n, m) \in \mathbb{Z} \times \mathbb{N}$; the *(homogeneous) approximation type* $L(x)$, in the terminology of J. C. Lagarias and J. O. Shallit [70], is the infimum of all $c > 0$ for which the same inequality admits at least one solution. Almost every real number has $L_\infty(x) = L(x) = 0$, by a classical result of A. Ya. Khinchin, but the numbers with $L_\infty(x) > 0$—the *badly approximable* numbers—form an uncountable class; Hurwitz's theorem gives $L_\infty(x) \le 1/\sqrt{5}$ for every irrational x, with equality at the golden ratio, and a theorem of A. Markov places $L_\infty(x) \in [0, 1/3]$ for every transcendental x. Determining these constants for specific numbers of interest is notoriously difficult: it is, for instance, a longstanding open problem whether $L_\infty(\pi)$ equals 0, or 1/3, or lies somewhere in between.

Now let

$$\mathcal{T} \;:=\; \left\{ \frac{\arg(p+qi)}{2\pi} \;:\; p, q \in \mathbb{Q},\; p^2 + q^2 = 1,\; pq \neq 0 \right\},$$

a countable dense subset of $(-\frac{1}{2}, \frac{1}{2})$ consisting exclusively of transcendental numbers: each element of $\mathcal{T}$ is the ratio of the logarithm of an algebraic number to $2\pi i$, and its transcendence follows from the Gelfond–Schneider theorem. The Lagrange constants and approximation types of the members of $\mathcal{T}$ are not currently known, and determining (or even approximating) these values for any particular $t \in \mathcal{T}$ is considered to be out of reach.

Theorem 51.

1. *If the Ultimate Positivity Problem is decidable for integer LRS of order* 6*, then the Lagrange constant $L_\infty(t)$ is computable for every $t \in \mathcal{T}$.*
2. *If the Positivity Problem is decidable for integer LRS of order* 6*, then the approximation type $L(t)$ is computable for every $t \in \mathcal{T}$.*

Here uniform computability means that there is a *single* algorithm that, given the rational point $p + qi$ determining $t \in \mathcal{T}$ together with a rational $\varepsilon > 0$, outputs a rational number within ε of the constant in question. The uniformity is an essential part of the statement: were the constants in question all, say, algebraic, then each would be computable individually—by an algorithm with the answer hard-wired—and the non-uniform statement would hold vacuously, regardless of the decidability hypothesis. Partial converses also hold, so that at order 6 even proofs of *undecidability* would have substantial Diophantine consequences; and analogous hardness results can be formulated in terms of *inhomogeneous* Diophantine approximation [99].

The sequences witnessing Theorem 51 are strikingly simple to describe. Given $t = \theta/2\pi \in \mathcal{T}$ and a positive rational r, consider

$$u_n \;=\; r\sin(n\theta) - n\big(1 - \cos(n\theta)\big), \qquad v_n \;=\; -r\sin(n\theta) - n\big(1 - \cos(n\theta)\big).$$

These are rational LRS of order 6, with characteristic roots 1, $e^{i\theta}$, and $e^{-i\theta}$, each of multiplicity two—the very configuration identified in Section 7.2 as the order-6 wall. The quantity $\max(u_n, v_n) = r\,|\sin(n\theta)| - n(1 - \cos(n\theta))$ is positive precisely when $n\theta$ lies exceptionally close to a multiple of 2π—at distance roughly $2r/n$—and a short calculation shows that $\langle -u_n \rangle$ and $\langle -v_n \rangle$ are both ultimately positive or not according to whether r/π lies below or above $L_\infty(t)$. Querying an Ultimate Positivity oracle for varying rational r therefore computes $L_\infty(t)$ to any desired precision; the argument for the approximation type, using a Positivity oracle on suffixes of the same sequences, is similar. The witnessing sequences are non-simple, as indeed they must be: simple LRS of order 6 fall within the scope of Theorem 50.

Simple LRS at order 10. For simple LRS, no Diophantine-hardness reduction of the above kind is known. What is known is that the techniques underlying Theorem 50 break down at order 10, and that one can exhibit concrete simple LRS of order 10 whose positivity no current technique appears able to settle; we present such a hard instance in Section 8.4. There is, moreover, the following reduction, showing that decidability of Positivity for simple LRS of order 10 would resolve the outstanding case of the Skolem Problem at order 5.

Theorem 52. *The Skolem Problem for LRS of order at most* 5 *Turing-reduces to the Positivity Problem for simple LRS of order* 10 *[95].*

The reduction appears in the thesis [95]; an observation to similar effect, with order 10 replaced by 14, was already made in [98]. It runs as follows. By the analysis of Section 5.3, the outstanding case of the Skolem Problem at order 5 concerns simple integer LRS of the form $u_n = w_n + b\,r^n$, where the dominant part w_n comprises the contributions of two complex-conjugate pairs of roots of common modulus ρ, and r is real with $|r| < \rho$. Fix a rational s with $r^2 < s < \rho^2$, and consider the sequence $\langle w_n^2 - s^n \rangle_{n=0}^{\infty}$: a simple LRS of order 10, ultimately positive by the growth theorem of the preceding subsection. Its *threshold*—an index N beyond which $w_n^2 > s^n$, whose existence the growth theorem guarantees only ineffectively—can now be *computed*, by querying the Positivity oracle on the successive suffixes $\langle w_{n+j}^2 - s^{\,n+j} \rangle_{n=0}^{\infty}$ for $j = 0, 1, 2, \ldots$ until the answer is affirmative. Beyond N (and beyond a further, easily computed index) we then have $w_n^2 > s^n > b^2 r^{2n}$, whence $u_n \neq 0$; the finitely many remaining indices are checked directly, and the Skolem Problem for $\boldsymbol{u}$ is thereby decided.

More generally, a Positivity oracle for simple LRS renders the ineffective growth theorem of Section 7.3 fully effective [63]. It is this effectivization that drives the algebraic-to-integer reduction of Theorem 47—which, we note, also disposes of any ambiguity in the field of definition of the oracle in Theorem 52—as well as the model-checking applications to which we return briefly in Section 9.

8 Worked Examples

We close the technical body of the survey with four worked examples, intended to illustrate the abstract material above on concrete sequences. The first two are pedagogical: the Fibonacci sequence, on which separation bounds and algebraic invariants can both be illustrated in miniature; and a small order-3 example on which Baker's Theorem, modular arguments, and p-adic techniques are all brought to bear. The remaining two are at the contemporary frontier: an order-6 sequence for which the Skolem Problem is currently open in spite of substantial computational evidence of zero-freeness, and an order-10 simple sequence for which the Positivity Problem is currently open and which falls outside the reach of every certificate class identified in Section 4.

8.1 Fibonacci, Revisited

We begin with the Fibonacci sequence $\boldsymbol{f} = \langle 0, 1, 1, 2, 3, 5, 8, 13, 21, \ldots \rangle$, satisfying $f_{n+2} = f_{n+1} + f_n$ with $f_0 = 0$ and $f_1 = 1$. Its characteristic polynomial $x^2 - x - 1$ has roots $\varphi = (1 + \sqrt{5})/2$ and $\tilde{\varphi} = (1 - \sqrt{5})/2$, and its closed form is

$$f_n \;=\; \frac{1}{\sqrt{5}}\varphi^n \;-\; \frac{1}{\sqrt{5}}\tilde{\varphi}^{\,n}.$$

The Fibonacci sequence is non-degenerate—the quotient $\varphi/\tilde{\varphi} = -\varphi^2$ has modulus greater than 1, and is therefore not a root of unity—and admits the unique zero $f_0 = 0$. The Skolem Problem for $\boldsymbol{f}$ is therefore settled by inspection.

Nevertheless, we can use this trivial-looking example to illustrate some of the key techniques that play a role in higher-order cases. The Archimedean separation bound for $|f_n|$ is elementary: since $\varphi^n > 1 > |\tilde{\varphi}|^n$ for all $n \geq 1$, we have

$$|f_n| \;\geq\; \frac{\varphi^n - |\tilde{\varphi}|^n}{\sqrt{5}} \;>\; 0 \qquad (n \geq 1),$$

which is a separation-bound certificate of zero-freeness on the tail $\{1, 2, 3, \ldots\}$. The bound is entirely straightforward; however, once the order grows and the dominant-root structure becomes more complex, such elementary estimates give way to the machinery of Baker's Theorem (Section 4.2).

We also remark that the bi-infinite completion of $\boldsymbol{f}$,

$$\langle \ldots, 5, -3, 2, -1, 1, 0, 1, 1, 2, 3, 5, \ldots \rangle,$$

contains the unique zero $f_0 = 0$, and any modular reduction $f_n \bmod m$ is purely periodic (and hence contains infinitely many zeros). As we remarked in Section 1, however, the *shifted* Fibonacci sequence $\langle 1, 1, 2, 3, 5, 8, \ldots \rangle$ provides the quintessential example of an LRS for which a modular argument fails to certify zero-freeness, despite the LRS having no zero on $\mathbb{N}$: every modular reduction contains the ghost of $f_0 = 0$.

The Fibonacci dynamical system and an algebraic invariant. In the matrix formulation of Section 3.2, the Fibonacci sequence arises by iterating the companion matrix $A := \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ on the initial vector $\boldsymbol{v}_0 = (1, 0)^\top$: the resulting orbit consists of the pairs (f_{n+1}, f_n) of consecutive Fibonacci numbers. Consider now the plane algebraic curve

$$I \;:\; (x^2 - xy - y^2 - 1)(x^2 - xy - y^2 + 1) \;=\; 0,$$

or, in expanded form, $x^4 + y^4 - 2x^3y - x^2y^2 + 2xy^3 - 1 = 0$: thus I is the union of the two hyperbolas $x^2 - xy - y^2 = \pm 1$, depicted in Figure 1. It is readily verified that the substitution $(x, y) \mapsto (x + y, x)$ takes the polynomial $x^2 - xy - y^2$ to $-(x^2 - xy - y^2)$, so that A maps each of the two hyperbolas onto the other; in particular, I is invariant under A, and since the initial point $(1, 0)$ lies on I, the entire orbit remains on I, alternating between the two hyperbolas (this is, in essence, Cassini's identity). We shall not explain here how I was obtained, nor prove that it is as small as possible; let us however recall from Section 4.4 that strongest algebraic invariants of linear dynamical systems are computable [60].

Imagine now that one wished to establish that the orbit never hits some given algebraic set $T \subseteq \mathbb{R}^2$. On the one hand, this is an instance of the Skolem Problem—of potentially higher order—since membership in T can be encoded as the vanishing of a single polynomial, which, evaluated along the orbit, yields an LRS [80]; it is equally the most basic instance of the reachability and model-checking problems for linear dynamical systems, to which we return in Section 9. If now T happens to be disjoint from I—a condition that can be mechanically decided by quantifier elimination over the reals—then the invariant I, together with the two verifications of the preceding paragraph, constitutes a certificate of non-reachability: that is, a certificate of zero-freeness for the corresponding Skolem instance. Likewise, if T is a semialgebraic (rather than algebraic) set, then, the matrix A being diagonalizable, non-reachability reduces to finitely many instances of the Positivity Problem [63]; and disjointness of T from I, again a decidable condition, once more furnishes a certificate.

8.2 An Order-3 Example: Baker at Work

For our second example, let $\boldsymbol{u}$ be the LRS defined by the exponential-polynomial formula

$$u_n = 25(4 + 3i)^n + 25(4 - 3i)^n - 24 \cdot 3^n.$$

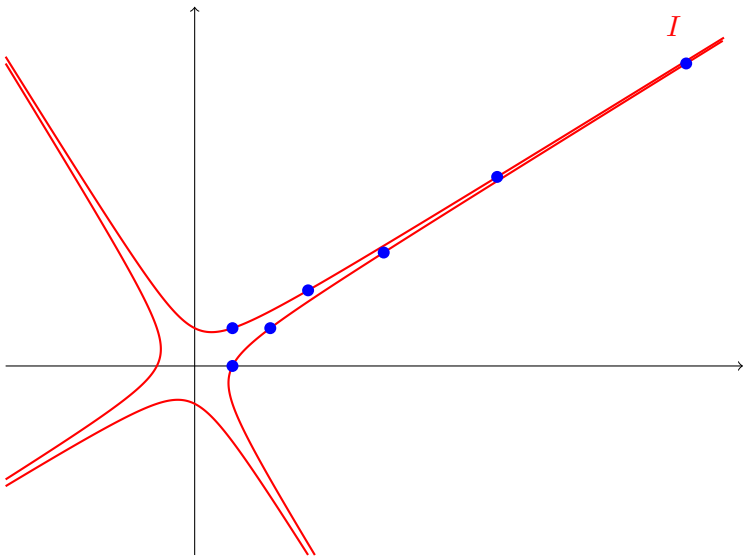

Figure 1: The orbit of $(1,0)$ under the Fibonacci companion matrix A (blue points: pairs of consecutive Fibonacci numbers) remains on the algebraic invariant I, the union of two hyperbolas that A maps onto one another.

Then $\boldsymbol{u}$ has three characteristic roots, $4+3i$, $4-3i$, and 3. They are all simple, and so $\boldsymbol{u}$ is simple. Moreover, one can verify that $\boldsymbol{u}$ is non-degenerate. The initial terms of $\boldsymbol{u}$ are $(u_0, u_1, u_2) = (26, 128, 134)$, and $\boldsymbol{u}$ obeys the linear recurrence

$$u_{n+3} = 11u_{n+2} - 49u_{n+1} + 75u_n.$$

Let us try to solve the Skolem Problem for $\boldsymbol{u}$ using a few different techniques. At the end, we also use $\boldsymbol{u}$ to highlight some aspects of the Positivity Problem.

Using Baker's Theorem. As $|4+3i| = |4-3i| = 5$ is larger than the third characteristic root 3, we have that $4+3i$ and $4-3i$ are dominant roots while 3 is a non-dominant root. Assume that $n \in \mathbb{N}$ satisfies $u_n = 0$, i.e., that $25(4+3i)^n + 25(4-3i)^n = 24 \cdot 3^n$. Dividing both sides by $|25(4-3i)^n|$ yields the equality

$$\left| -\left(\frac{4+3i}{4-3i}\right)^n - 1 \right| = \frac{24}{25}\left(\frac{3}{5}\right)^n, \tag{12}$$

which we now connect to Theorem 7. Set $\lambda := \frac{4+3i}{4-3i}$, an algebraic number of modulus 1 which, by non-degeneracy, is not a root of unity. We may assume that the left-hand side of (12) is at most $\frac{1}{2}$, so that λ^n lies close to -1; choose $k \in \mathbb{Z}$ with $\left|n \arg \lambda - (2k+1)\pi\right| \leq \pi$, noting that $|2k+1| \leq n+1$, and consider the linear form in logarithms

$$\Lambda \;:=\; n \log \lambda \;-\; (2k+1)\log(-1),$$

where log denotes the principal branch, so that $\log \lambda = i \arg \lambda$ and $\log(-1) = i\pi$. Since λ is not a root of unity we have $\lambda^n \neq -1$, hence $\Lambda \neq 0$, and Theorem 7—applied to the two logarithms $\log \lambda$ and $\log(-1)$ with integer coefficients n and $-(2k+1)$—yields $|\Lambda| \geq (n+1)^{-C}$ for an effectively computable constant $C > 0$. On the other hand, the elementary inequality $|e^{iy} - 1| = 2\,|\sin(y/2)| \geq 2|y|/\pi$, valid for $|y| \leq \pi$, gives

$$\left| -\left(\frac{4+3i}{4-3i}\right)^n - 1 \right| \;=\; |e^{\Lambda} - 1| \;\geq\; \frac{2}{\pi}\,|\Lambda|.$$

Combining the above with (12), we see that $u_n = 0$ implies

$$\frac{24}{25}\left(\frac{3}{5}\right)^n \geq \frac{2}{\pi}(n+1)^{-C}.$$

This inequality holds only for n below an effectively computable threshold; testing whether $u_n = 0$ for each n below that threshold then decides the Skolem Problem for $\boldsymbol{u}$.

Using p-adic Baker. Besides the usual Archimedean variant of Baker's Theorem, we can also employ the $\mathfrak{p}$-adic version for a prime ideal $\mathfrak{p}$ of $\mathbb{Z}[i]$. For almost all such prime ideals $\mathfrak{p}$, for example (7) or $(4+i)$, all three characteristic roots are dominant with respect to the absolute value $|\cdot|_\mathfrak{p}$. In these cases, we cannot apply a $\mathfrak{p}$-adic version of Baker's Theorem. The only exceptions are the prime ideals above 3 and 5, the rational primes dividing the product 75 of the characteristic roots. In $\mathbb{Z}[i]$, $\mathfrak{p} := (3)$ is a prime ideal, $4+3i$ and $4-3i$ are $\mathfrak{p}$-dominant and not in $\mathfrak{p}$, and the last characteristic root, 3, is in $\mathfrak{p}$ and is therefore not $\mathfrak{p}$-dominant.

Let $n \in \mathbb{N}$ be such that $u_n = 0$. Then, as before, we have that $25(4+3i)^n + 25(4-3i)^n = 24 \cdot 3^n$, whence both sides have the same 3-adic valuation:

$$v_3\big(25(4+3i)^n + 25(4-3i)^n\big) = v_3(24 \cdot 3^n) = 1+n.$$

Dividing once more by $-25(4-3i)^n$, which has 3-valuation 0,[4] we conclude that $v_3(-\lambda^n - 1) = 1+n$, where, as before, $\lambda = \frac{4+3i}{4-3i}$. On the other hand, the $\mathfrak{p}$-adic version of Baker's Theorem (Theorem 8)—applied to the two algebraic numbers -1 and λ with respective exponents 1 and n, the quantity $-\lambda^n - 1$ being non-zero since λ is not a root of unity—yields an effectively computable constant $C > 0$ such that, for all $n \geq 3$,

$$v_3\big(-\lambda^n - 1\big) \;<\; C \log n.$$

Hence, $u_n = 0$ would imply that $1 + n < C \log n$, which again effectively bounds n. Thus, with this method we can also decide the Skolem Problem for $\boldsymbol{u}$.

Using modular arguments. Compared to the previous lines of reasoning which relied on deep number-theoretical results, there is a far easier argument: we observe that $u_n \equiv 2 \pmod 3$ and $u_n \not\equiv 0 \pmod 5$ for all $n \in \mathbb{N}$. Either of these observations is sufficient to certify that $u_n \neq 0$ for all $n \in \mathbb{N}$. However, for any other prime p, there are infinitely many $n \in \mathbb{N}$ such that $u_n \equiv 0 \pmod p$ as $u_{-1} = 0$ and the recurrence can be run backward modulo p. Also note that the observations above only apply for $n \in \mathbb{N}$: we have that $u_{-1} = 0$, and u_n modulo 3 or 5 is not even defined for $n < -1$. There are many LRS for which Baker-style methods apply while such an easy modular argument does not exist.

We could also use the more general technique of Theorem 33, as implemented in the SKOLEM tool, since $\boldsymbol{u}$ is simple. Fixing the prime $p = 2$ and studying the 2-adic interpolation, one can show that $v_2(u_{4n-1}) = v_2(n) + 5$ for all $n \in \mathbb{Z}$, so that $u_{4n-1} \neq 0$ for all $n \in \mathbb{Z} \setminus \{0\}$, and similarly that $v_2(u_{4n}) = 1$ for all $n \in \mathbb{Z}$, so that $u_{4n} \neq 0$ throughout. For the remaining indices, $n \equiv 1, 2 \pmod 4$, one checks that $u_n \not\equiv 0 \pmod{13}$. Altogether, $u_n \neq 0$ for all $n \in \mathbb{Z} \setminus \{-1\}$, and in particular for all $n \in \mathbb{N}$.

Positivity. Now let us study the Positivity Problem for $\boldsymbol{u}$, which can be handled quite easily. Writing $4+3i = 5e^{i\theta}$, with θ/π irrational (a consequence of non-degeneracy), we obtain that

$$u_n = 50 \cdot 5^n \cos(n\theta) - 24 \cdot 3^n.$$

As θ/π is irrational, the set $\{\cos(\theta n) : n \in \mathbb{N}\}$ is dense in $[-1, 1]$. Therefore, u_n will be negative for infinitely many n.

[4] Here the valuation function v_3 is considered over the domain $\mathbb{Q}(i)$.

Hence, for the sake of exposition, let us study the LRS $\boldsymbol{w}$ defined by $w_n = u_n + c \cdot 5^n$ for a fixed rational number c. First, if $c < 50$, then by the exact same argument, w_n will be negative for infinitely many $n \in \mathbb{N}$. If $c > 50$, we have that

$$w_n = 25(4+3i)^n + 25(4-3i)^n + c \cdot 5^n - 24 \cdot 3^n \geq (c-50)5^n - 24 \cdot 3^n,$$

and so one can easily find a bound N such that $w_n \geq 0$ for all $n \geq N$, leaving finitely many initial indices to check directly.

Thus, only the case $c = 50$ remains, for which we need Baker's Theorem. In that case, we have two parts of the LRS 'racing' each other: the positive part coming from the dominant roots, $25((4+3i)^n + (4-3i)^n + 2 \cdot 5^n)$, and the negative part coming from the non-dominant root, $-24 \cdot 3^n$. We therefore have to estimate how small $(4+3i)^n + (4-3i)^n + 2 \cdot 5^n$ can become in comparison to $24 \cdot 3^n$. We observe that, choosing the square roots so that their product is 5,

$$(4+3i)^n + (4-3i)^n + 2 \cdot 5^n = \left(\sqrt{4+3i}^{\,n} + \sqrt{4-3i}^{\,n}\right)^2.$$

Thanks to Baker's Theorem, the term inside the square can be lower-bounded in absolute value by $n^{-C}\sqrt{5}^n$ for some effectively computable constant $C > 0$. Squaring, we find that $25n^{-2C}5^n$ eventually exceeds $24 \cdot 3^n$, with an effectively computable threshold, and so we can decide Positivity for the LRS $\boldsymbol{w}$.

8.3 An Order-6 Sequence at the Skolem Frontier

Our third example is the order-6 LRS

$$u_n \;=\; 2(-4+7i)^n + 2(-4-7i)^n + 4(8+i)^n + 4(8-i)^n + n, \tag{13}$$

which we first met in Section 4.3, and to which Sections 5.4 and 6.3 have promised a return. The characteristic polynomial of this LRS is

$$p(x) \;=\; (x-1)^2(x^2+8x+65)(x^2-16x+65),$$

with the double root at 1 contributing the linear-in-n term, and with four simple roots $-4 \pm 7i$ and $8 \pm i$: Gaussian integers of norm $65 = 5 \cdot 13$, hence of common Archimedean modulus $\sqrt{65}$. One can verify that $\boldsymbol{u}$ is non-degenerate. The first few values are

$$u_0 = 12, \quad u_1 = 49, \quad u_2 = 374, \quad u_3 = 6003, \quad u_4 = 21520, \quad u_5 = 150773,$$

and the sequence takes each sign infinitely often (the first negative value being u_{11}), ruling out any analysis based on eventual sign control of a dominant term.

Outside the MSTV class. The sequence (13) satisfies neither condition of Definition 22. It has four dominant roots in modulus, namely $-4 \pm 7i$ and $8 \pm i$: one too many for the first condition. As for the second, at a non-Archimedean place v not lying above 5 or 13, all five distinct characteristic roots are v-adic units, and hence all five are v-dominant; whereas at each of the four places above 5 and 13, exactly two of the Gaussian roots acquire absolute value less than 1—each of $-4 \pm 7i$, $8 \pm i$ is divisible by exactly one of the two prime ideals above 5, and likewise above 13—leaving three v-dominant roots, namely two Gaussian roots together with the root 1. No non-Archimedean place therefore has at most two dominant roots, and the second condition fails as well. Theorem 23 does not apply.

Three certificate classes, three failures. We now bring the certificate framework of Section 4 to bear on sequence (13). None of the three classes settles the Skolem question—and for two of the three, one can *prove* that no certificate exists.

Separation-bound certificates. Both the Archimedean and p-adic separation methods of Section 4.2 founder on the dominant-root counts just established: at every place, at least three dominant roots may in principle combine destructively at sporadic indices, and Baker's Theorem in its current form cannot rule this out.

Modular invariants. No modular certificate exists at all: as Proposition 53 below shows, the sequence has zeros modulo every modulus $m \geq 2$.

Semialgebraic invariants. By Theorem 19, the existence of a semialgebraic invariant for this LRS—a closed semialgebraic set $I \subseteq \mathbb{R}^6$ containing the initial state, closed under the companion-matrix dynamics, and disjoint from the hyperplane $\{x_1 = 0\}$—is decidable; running the algorithm of [9] on this particular sequence shows that no such invariant exists.

Proposition 53. *For every integer $m \geq 2$ there exist infinitely many $n \in \mathbb{N}$ such that $u_n \equiv 0 \pmod{m}$.*

Proof. Write $u_n = V_n + n$, where

$$V_n \;:=\; 2(-4+7i)^n + 2(-4-7i)^n + 4(8+i)^n + 4(8-i)^n$$

is a simple integer LRS whose characteristic roots are Gaussian integers of norm 65. The linear term n will play the role of a unit derivative in a Hensel-style lifting argument, while the exponential part $\boldsymbol{V}$ is kept under control by the following periodicity property.

Step 1: periodicity. For every prime p and every $E \geq 1$ there is a positive integer $T = T(p, E)$ with $v_p(T) \leq E - 1$ such that

$$V_{n+T} \;\equiv\; V_n \pmod{p^E} \qquad \text{for all } n \geq E. \tag{14}$$

Suppose first that p is odd and $p \notin \{5, 13\}$, and set $T := (p^2 - 1)\, p^{E-1}$. Every characteristic root α of $\boldsymbol{V}$ is a unit modulo each prime ideal $\mathfrak{p}$ of $\mathbb{Z}[i]$ above p; since the residue field has at most p^2 elements, $\alpha^{p^2-1} \equiv 1 \pmod{\mathfrak{p}}$, and raising to the pth power $E-1$ times gives $\alpha^T \equiv 1 \pmod{\mathfrak{p}^E}$ (if $x \equiv 1 \pmod{\mathfrak{p}^c}$ with $c \geq 1$, then $x^p \equiv 1 \pmod{\mathfrak{p}^{c+1}}$, as p is odd and unramified in $\mathbb{Z}[i]$). Hence $V_{n+T} - V_n = \sum_j c_j\, \alpha_j^n (\alpha_j^T - 1) \equiv 0 \pmod{p^E}$, in fact for all $n \geq 0$. Next let $p \in \{5, 13\}$, again with $T := (p^2 - 1)\, p^{E-1}$. Now $p = \mathfrak{p}\overline{\mathfrak{p}}$ splits in $\mathbb{Z}[i]$, and each of the four roots is divisible by exactly one of $\mathfrak{p}$, $\overline{\mathfrak{p}}$. Modulo $\mathfrak{p}^E$ the terms of the two $\mathfrak{p}$-divisible roots are congruent to 0 as soon as $n \geq E$, while the other two roots are $\mathfrak{p}$-units, to which the previous reasoning applies; arguing symmetrically at $\overline{\mathfrak{p}}$ and combining the two congruences yields (14) for $n \geq E$. Finally let $p = 2$. Since the two roots of each conjugate pair carry equal coefficients, $V_n = 4\,\mathrm{Re}\big((-4+7i)^n\big) + 8\,\mathrm{Re}\big((8+i)^n\big) \equiv 0 \pmod 4$, so for $E \leq 2$ we may take $T := 1$. For $E \geq 3$, every element of $\mathbb{Z}[i]$ of odd norm satisfies $\alpha^4 \equiv 1 \pmod 8$ (the unit group of $\mathbb{Z}[i]/8\mathbb{Z}[i]$ has exponent 4), and repeated squaring gives $\alpha^{2^{E-1}} \equiv 1 \pmod{2^E}$; thus $T := 2^{E-1}$ works, again for all $n \geq 0$.

Step 2: lifting. Let $W = W' p^e$ with p prime, $p \nmid W'$ and $e \geq 0$, and suppose that $u_{n_0} \equiv 0 \pmod W$. Suppose further that S is a positive integer such that: (i) $V_{n+S} \equiv V_n \pmod{W' p^{e+1}}$ for all $n \geq n_0$; (ii) $W \mid S$; and (iii) $v_p(S) = e$. Then for every $\lambda \geq 0$,

$$u_{n_0+\lambda S} \;=\; V_{n_0+\lambda S} + n_0 + \lambda S \;\equiv\; u_{n_0} + \lambda S \pmod{W' p^{e+1}},$$

and since $S p^{-e}$ is invertible modulo p, an appropriate choice of λ—which may moreover be taken arbitrarily large, as only its residue modulo p matters—gives $u_{n_0+\lambda S} \equiv 0 \pmod{W' p^{e+1}}$, the congruence modulo W' being preserved thanks to (ii). The solution thus lifts from the modulus W to the modulus $W' p^{e+1}$.

Step 3: exhaustion. Let $p_1 < p_2 < \cdots$ enumerate the primes, and set $M_k := \prod_{j=1}^{k} p_j^{3k}$. Since every $m \geq 2$ divides M_k for all sufficiently large k, it suffices to solve $u_n \equiv 0 \pmod{M_k}$ for every k, with n arbitrarily large. The base case is the direct computation $u_4 = 21520 \equiv 0 \pmod{2^3}$. For the inductive step, suppose $u_{n_k} \equiv 0 \pmod{M_k}$. Applying Step 2 repeatedly, we raise the exponent of each of $p_1, \ldots, p_k$ from $3k$ to $3k+3$, one prime and one exponent at a time, and then admit the new prime p_{k+1} and raise its exponent from 0 to $3k+3$ likewise, arriving at M_{k+1}. Crucially, in each application of Step 2 one may take for S simply the current modulus W: condition (ii) is then trivial; condition (iii) holds because $v_p(W) = e$; and condition (i) reduces, by the Chinese Remainder Theorem and Step 1, to the divisibilities $T(r, F) \mid W$ for the prime powers r^F exactly dividing the target modulus $W' p^{e+1}$. These divisibilities in turn follow from two elementary observations: every prime factor of $r^2 - 1$, for odd prime r, is smaller than r (the odd part of $r+1$ is at most $(r+1)/2$); and the exponents in the schedule M_k grow fast enough that $v_s(r^2-1) \leq 3k$ for all primes $s < r \leq p_{k+1}$, as one checks using the crude bound $p_{k+1} < 2^{2k}$, itself an easy consequence of Bertrand's postulate. Finally, taking λ large in each application of Step 2 keeps n beyond the threshold $n \geq E$ required by Step 1, and produces solutions that are arbitrarily large; in particular, for each fixed m the solution set is infinite. □

The zeros whose existence Proposition 53 guarantees are thus all, for all we currently know, ghosts in the sense of Section 1: none of them need lift to an integer zero of $\boldsymbol{u}$.

What is known: a 17-adic analysis. The most substantial positive information about sequence (13) comes from the p-adic machinery of Sections 5.4 and 6.3, which we now see in action. The prime $p = 17$ suits the sequence well: since $4^2 \equiv -1 \pmod{17}$, the prime 17 splits in $\mathbb{Q}(i)$, and the embedding of $\mathbb{Z}[i]$ into $\mathbb{Z}_{17}$ sending i to the square root of -1 congruent to 4 modulo 17 places all five characteristic roots in $\mathbb{Z}_{17}$, with reductions

$$-4+7i \;\mapsto\; 7, \qquad -4-7i \;\mapsto\; 2, \qquad 8+i \;\mapsto\; 12, \qquad 8-i \;\mapsto\; 4, \qquad 1 \;\mapsto\; 1 \pmod{17}.$$

All five residues are non-zero, so all roots are 17-adic units, and their multiplicative orders modulo 17—respectively 16, 8, 16, 4, and 1—have least common multiple $L = 16$. Exactly as in the proof of the Skolem–Mahler–Lech Theorem (Section 5.2), each of the sixteen subsequences $\langle u_{16n+r} \rangle_{n=0}^{\infty}$, for $0 \leq r < 16$, is therefore interpolated by a 17-adic analytic function $f_r : \mathbb{Z}_{17} \to \mathbb{Z}_{17}$.

Running the residue-class search of Section 6.3 (as implemented in the p-adic mode of the SKOLEM tool [31]) on these sixteen interpolants terminates within approximately 45 seconds and returns a complete inventory: the sequence $\boldsymbol{u}$ has exactly sixteen 17-adic zeros, one in each of the sixteen subsequences $\langle u_{16n+r} \rangle$, each of multiplicity 1, each isolated in its own disc by Hensel's Lemma, with the Newton-polygon zero counts certifying that none has been missed. The base-17 expansion of each zero can moreover be computed to any desired depth. In particular, each of the sixteen zeros turns out to have a non-zero base-17 digit at position 818, whereas a natural number $n \leq 10^{1000} < 17^{816}$ has all base-17 digits at positions 816 and beyond equal to zero. None of the sixteen 17-adic zeros is therefore a natural number below 10^{1000}; and as every natural-number zero of $\boldsymbol{u}$ is in particular a 17-adic zero, we conclude—this is the certification promised in Section 5.4—that

$$u_n \;\neq\; 0 \qquad \text{for all } n \in \{0, 1, \ldots, 10^{1000}\}.$$

An elementary supplement to the tool's output localizes any integer zero rather precisely. First, since each conjugate pair of dominant roots carries equal coefficients, we have $V_n \equiv 0 \pmod 4$, where $V_n := u_n - n$ as in the proof of Proposition 53; thus $u_n \equiv n \pmod 4$, and any zero must satisfy $n \equiv 0 \pmod 4$. Next, writing $n = 4m$ and recalling from the same proof that every element of $\mathbb{Z}[i]$ of odd norm satisfies $\alpha^4 \equiv 1 \pmod 8$, each conjugate pair contributes $\alpha^{4m} + \overline{\alpha}^{4m} \equiv 2 \pmod{16}$ (indeed, writing $\alpha^{4m} = 1 + 8\gamma$ we have $\alpha^{4m} + \overline{\alpha}^{4m} = 2 + 16\,\mathrm{Re}\,\gamma$), whence

$$u_{4m} \;\equiv\; 2\,(2+4) + 4m \;\equiv\; 12 + 4m \pmod{16},$$

and $u_{4m} \equiv 0 \pmod{16}$ forces $m \equiv 1 \pmod 4$, that is, $n \equiv 4 \pmod{16}$. Consequently, of the sixteen subsequences, only $\langle u_{16n+4} \rangle$ can contain an integer zero; and since its interpolant f_4 accounts for exactly one of the sixteen 17-adic zeros, the sequence (13) has *at most one* zero $n \in \mathbb{N}$—necessarily congruent to 4 modulo 16 and, by the digit analysis above, larger than 10^{1000}.

What this analysis cannot deliver is a verdict on the unbounded Skolem Problem: the single remaining candidate—the unique 17-adic zero of f_4—is, as far as anyone knows, a natural number (necessarily astronomically large), some other rational number, or an irrational element of $\mathbb{Z}_{17}$; and, as discussed in Section 6.3, no effective procedure is known that can tell these alternatives apart. To settle the Skolem Problem for sequence (13), one would need an *a priori* effective bound on the magnitude of its integer zeros—precisely what Baker-type methods provide inside the MSTV class, and what is missing outside it.

An aside: when minor modifications restore tractability. It is instructive to compare sequence (13) with the slight modification

$$\widetilde{u}_n \;:=\; 2(-4+7i)^n + 2(-4-7i)^n \;+\; 2(8+i)^n + 2(8-i)^n + n, \tag{15}$$

in which the four dominant roots now carry the *same* coefficient 2. Writing $\theta_1 := \arg(-4+7i)$ and $\theta_2 := \arg(8+i)$, the dominant part of $\widetilde{u}_n$ collapses, via the sum-to-product identity for cosines, into

$$4\sqrt{65}^{\,n}\big(\cos n\theta_1 + \cos n\theta_2\big) \;=\; 8\sqrt{65}^{\,n}\,\cos\Big(n\,\tfrac{\theta_1+\theta_2}{2}\Big)\cos\Big(n\,\tfrac{\theta_1-\theta_2}{2}\Big),$$

a product of two cosines, each of which can be bounded away from zero by Baker's Theorem, much as in Section 8.2. This yields an effective bound on the magnitude of any zero of $\widetilde{\boldsymbol{u}}$, and hence decidability of its Skolem Problem; an argument of exactly this shape is carried out in detail in [31, Appendix A]. A subsequent computation with the SKOLEM tool does indeed verify that $\widetilde{u}_n \neq 0$ for all $n \in \mathbb{N}$. Since $\widetilde{\boldsymbol{u}}$ has the same characteristic roots as $\boldsymbol{u}$, it lies equally outside the MSTV class: membership in the class is thus sufficient, but by no means necessary, for effective zero bounds. The genuine obstacle in sequence (13) is the *asymmetry* of the coefficients 2 and 4: the quantity $\cos n\theta_1 + 2\cos n\theta_2$ admits no comparable factorization, and bounding it away from zero would require lower bounds on linear forms of a shape that present-day methods do not deliver. What new certificate technique might suffice to handle such asymmetric configurations is, in our view, one of the central open questions in this area.

8.4 An Order-10 Sequence at the Positivity Frontier

Our final example concerns the Positivity Problem. Consider the simple order-10 integer LRS

$$u_n \;:=\; w_n^2 - 2^n, \tag{16}$$

where

$$w_n \;:=\; \big((1+2i)(2+3i)\big)^n + \big((1-2i)(2-3i)\big)^n \;+\; 2\big((1-2i)(2+3i)\big)^n \;+\; 2\big((1+2i)(2-3i)\big)^n.$$

Equivalently, $\boldsymbol{u}$ satisfies the order-10 recurrence

$$\begin{aligned} u_{n+10} \;=\;& -u_{n+9} - 378u_{n+8} + 749576u_{n+7} - 2333386u_{n+6} + 55996590u_{n+5} \\ & - 205750047100u_{n+4} + 856834394000u_{n+3} + 13815580471875u_{n+2} \\ & + 20682499470546875u_{n+1} - 41423825675781250u_n, \end{aligned}$$

with initial values $u_0, u_1, \ldots, u_9$ given by 35, 574, 34592, 8999992, 115734548, 5682747424, 1837938758372, 13061285121472, 397924220049188, 290333397927490624.

The construction. The building blocks are the four Gaussian primes $1\pm2i$ and $2\pm3i$. Setting $\lambda_1 := (1+2i)(2+3i) = -4+7i$ and $\lambda_2 := (1+2i)(2-3i) = 8+i$, the four products of one prime from each conjugate pair are $\lambda_1, \overline{\lambda_1}, \lambda_2, \overline{\lambda_2}$—precisely the dominant roots of sequence (13) from the previous section, of common modulus $\sqrt{65}$. The connection is no accident: $2w_n$ is exactly the exponential part of sequence (13), so our two frontier examples are driven by the same exponential sum, with the same asymmetric coefficient pattern. Squaring w_n produces a simple LRS of order at most $\binom{4+1}{2} = 10$, by the Hadamard-square analysis of Section 3.4: its characteristic roots are the pairwise products of the roots of $\boldsymbol{w}$, namely the eight non-real Gaussian integers

$$\lambda_1^2 = -33-56i, \quad \lambda_2^2 = 63+16i, \quad \lambda_1\lambda_2 = -39+52i, \quad \lambda_1\overline{\lambda_2} = -25+60i$$

and their conjugates, together with $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2} = 65$: nine distinct roots, all of modulus 65. The subtracted exponential 2^n contributes the tenth root 2; and as all ten coefficients of the resulting exponential-polynomial representation are non-zero, $\boldsymbol{u}$ is simple of order exactly 10, with nine dominant roots and the single non-dominant root 2. One checks as usual that $\boldsymbol{u}$ is non-degenerate.

The same recipe yields an entire family of hard instances: take two Gaussian primes p_1, p_2 of distinct odd norms, neither real nor purely imaginary; set $\lambda_1 := p_1p_2$ and $\lambda_2 := p_1\overline{p_2}$; attach non-zero Gaussian coefficients a_1, a_2 of distinct moduli, forming $w_n := a_1\lambda_1^n + \overline{a_1}\,\overline{\lambda_1}^n + a_2\lambda_2^n + \overline{a_2}\,\overline{\lambda_2}^n$; and set $u_n := w_n^2 + a\,r^n$, where r is a rational integer with $2 \le |r| < |\lambda_1|^2$ and a is a non-zero rational integer, at least one of a and r being negative. Sequence (16) is the instance $p_1 = 1+2i$, $p_2 = 2+3i$, $(a_1, a_2) = (1, 2)$, $(a, r) = (-1, 2)$.

Ultimately positive; Positivity open. Since $\boldsymbol{w}$ is simple and non-degenerate, with all four roots of modulus $\sqrt{65}$, the growth theorem encountered in Section 7.3 [53, 123] gives $|w_n| > \sqrt{2}^n$ for all sufficiently large n; hence $u_n = w_n^2 - 2^n > 0$ eventually, and sequence (16) is ultimately positive. (The same argument applies to every instance of the general scheme, thanks to the requirement $|r| < |\lambda_1|^2$.) But the growth theorem is ineffective: as discussed in Section 7.3, its known proofs rest on the Subspace Theorem and yield no bound whatsoever on the onset index. Nor is this ineffectivity incidental. The sequence (16) belongs to exactly the family $\langle w_n^2 - s^n\rangle$—here with $s = 2$—through which the outstanding order-5 case of the Skolem Problem Turing-reduces to Positivity for simple LRS of order 10 (Theorem 52): the ability to compute onset indices for instances of this shape would resolve the Skolem Problem at order 5. Conversely, an oracle for simple Positivity would render the onset effective, as explained in Section 7.4 [63].

What is known. Direct computation confirms that $u_n \ge 0$ for all $0 \le n \le 10^6$. Moreover, the SKOLEM tool verifies that $\boldsymbol{u}$ has no zeros at all; in notable contrast to sequence (13), the *Skolem* Problem for $\boldsymbol{u}$ is thus resolved, while *Positivity* remains open. Between the computed initial range and the ineffective onset of the growth theorem lies a gap of indeterminate length, and whether some index in that gap has $u_n < 0$ is the open Positivity question on $\boldsymbol{u}$. In addition, for any $\ell \in \mathbb{Z}$, the SKOLEM tool can in principle compute all indices n for which $u_n = \ell$, since the corresponding LRS remains simple.

The certificate landscape, again. Modular and p-adic certificates being of no use for Positivity (Section 4.5), the two remaining classes are Archimedean separation bounds and semialgebraic invariants, and neither currently suffices on sequence (16).

Separation bounds. Closing the gap would require an effective lower bound on $|w_n|$ against $\sqrt{2}^n$: a Baker-type analysis of the four dominant Gaussian roots of $\boldsymbol{w}$. But these roots carry the same asymmetric coefficients $1, 1, 2, 2$ that placed sequence (13) beyond the reach of present-day methods: no trigonometric collapse is available, and no suitable extension of Baker's Theorem is known.

Semialgebraic invariants. By Theorem 19, existence of a semialgebraic invariant for the sequence (16) certifying $u_n \geq 0$ is decidable. The algorithm of [9] again returns a negative verdict: no such invariant exists.

Why this matters. The sequence (16) sits at the boundary of what is currently decidable on the Positivity side; in the spirit of Section 4.6, it provides a definite target against which to measure any prospective new certificate technique.

9 Related Problems and Frontiers

We close the survey with a tour of the broader landscape surrounding the Skolem and Positivity Problems: cognate problems that arise in different mathematical settings, application areas in which Skolem and Positivity feature prominently as algorithmic benchmarks, and the principal open problems that we believe will shape the next decade of research.

9.1 The Skolem Problem in Other Settings

Throughout the survey we have restricted attention to LRS over fields of characteristic 0. Several other natural settings have been studied, and we collect here pointers to the relevant literature.

Positive characteristic. Over fields of positive characteristic, the structure of the zero set of an LRS is considerably richer than the union-of-arithmetic-progressions characterization of the Skolem–Mahler–Lech Theorem. The foundational result is due to H. Derksen [49], who showed that the zero set of an LRS over a field of characteristic p is a so-called *p-normal* set, definable in terms of p-automatic languages. Crucially, Derksen's result is fully effective: he gives an algorithm that, given an LRS over $\overline{\mathbb{F}_p}(t)$, computes a finite-automaton description of its zero set, thereby resolving the Skolem Problem in this setting. Subsequent work of H. Derksen and D. Masser [50] extends the analysis to multivariate exponential-polynomial equations over fields of positive characteristic, giving effective bounds on the structure of the solution set: in particular, given LRS $u_1, \ldots, u_m$ over such a field, the set of tuples $(k_1, \ldots, k_m)$ with $u_1(k_1) + \cdots + u_m(k_m) = 0$ is effectively computable. The same paper shows, by way of contrast, that in characteristic 0 the existence of such a tuple is *undecidable* for LRS over the algebraic numbers,[5] confirming a conjecture of L. Cerlienco, M. Mignotte, and F. Piras. Most recently, R. Dong and O. Shafrir [51] have established the decidability of the Skolem Problem over arbitrary finitely generated commutative rings of positive characteristic, building on Derksen's framework together with a recent result on solving linear equations over powers of multiplicatively independent numbers.

Continuous-time analogues. The Skolem and Positivity Problems have natural analogues in the continuous-time setting, where the discrete sequence $\boldsymbol{u} = \langle u_n \rangle$ is replaced by a smooth function $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ satisfying a homogeneous linear differential equation with constant coefficients (equivalently, an exponential-polynomial function over $\mathbb{R}$). The Continuous Skolem Problem then asks whether such an f has a zero in $\mathbb{R}_{\geq 0}$, and the Continuous Positivity Problem asks whether f is non-negative on $\mathbb{R}_{\geq 0}$. The continuous setting differs substantially from the discrete one in both technique and typical results: the p-adic methods central to the discrete case are unavailable, and one relies instead on real-analytic techniques and on results from o-minimality and model theory. Decidability and hardness results in the continuous setting appear in [23, 44, 42, 45], among others. Let us finally mention [85], in which the authors extend the

[5] Derksen and Masser's undecidability proof makes use of $m = 557844$ distinct LRS, although that number could no doubt be drastically reduced.

algorithm synthesizing strongest algebraic invariants of [59] to the class of *linear hybrid automata*, which comprise both discrete and continuous components.

Robust and pseudo-variants. A natural strengthening of the Skolem and Positivity Problems asks for these properties to hold not just for the given LRS, but for every LRS in a small neighborhood—a robust analogue motivated by the verification of systems with imperfect or noisy parameters. The Pseudo-Skolem Problem is decidable [47], and the related Pseudo-Reachability Problem is decidable for diagonalizable linear dynamical systems [48]. A complementary set of robust-Skolem and robust-Positivity results is established in [5]. The pattern across these results is that robust analogues are typically *easier* than their exact counterparts, since small perturbations smooth out the delicate Diophantine phenomena that make the exact problems hard.

Rounding. A different relaxation considers linear dynamical systems in which the state is *rounded* after each application of the matrix, modeling finite-precision implementations of linear loops. With rounding to a fixed granularity, point-to-point reachability is PSPACE-complete for hyperbolic systems (those whose matrix has no eigenvalue of modulus 1), and decidability extends to arbitrary spectra for several natural classes of rounding functions [16]. Under floating-point rounding [74], the orbits of systems with non-negative matrices are eventually structured—the mantissas are periodic and the exponents grow linearly—and ω-regular model checking against semialgebraic predicates becomes decidable, in contrast with the unrounded setting, where the non-negative case already encompasses the Skolem and Positivity Problems. With negative matrix entries, however, reachability becomes undecidable, whereas exact point-to-point reachability is decidable in polynomial time [61].

9.2 Skolem and Positivity as Benchmarks

The Skolem and Positivity Problems serve as central reference points for hardness in a wide range of decision questions arising across theoretical computer science. We have already mentioned several such areas in the introduction; we expand here on a representative selection. The list below is far from exhaustive and reflects our particular interests.

Reachability for linear dynamical systems. Several classical decision problems for discrete linear dynamical systems—systems evolving by repeated application of a fixed rational matrix A to an initial point $\boldsymbol{x}$—reduce to (or from) Skolem and Positivity. The Orbit Problem—given A and points $\boldsymbol{x}, \boldsymbol{y}$, does $A^n\boldsymbol{x} = \boldsymbol{y}$ for some n?—is decidable in polynomial time [61], but its parametric extensions encounter Skolem-style obstacles [15]. When the target is a vector subspace rather than a point, decidability is known for targets of dimension at most 3 [43] and, more recently, for targets of dimension logarithmic in a suitably defined dimension of the orbit, via the Simultaneous Skolem Problem of Section 6.3; targets of dimension linear in that of the orbit, on the other hand, are as hard as the Skolem Problem [14]. The Polyhedron-Hitting Problem—does the orbit of $\boldsymbol{x}$ ever enter a given polyhedron?—is decidable in restricted dimensions but Skolem-hard in general [41]. The algebraic and semialgebraic model-checking frameworks of [80, 63] cast a broad family of such reachability questions (and beyond) in a unified setting and locate each relative to Skolem and Positivity hardness. For a comprehensive treatment of the model-checking problem for linear dynamical systems, we refer the reader to the doctoral thesis of T. Karimov [62]; see also [64].

Loop termination and invariant synthesis. The question of whether a simple linear loop over the integers terminates is intimately connected to the Skolem Problem [101, 65]. On the

invariant side, the strongest algebraic invariant of a linear loop is computable, as noted in Section 4.4, while the synthesis of strongest algebraic invariants for single-path polynomial programs is Skolem-hard [92].

Markov chains and probabilistic systems. Quantitative verification of probabilistic systems is a particularly rich source of Skolem- and Positivity-hard problems: reachability for Markov chains [2, 122], the model checking of Markov chains viewed as distribution transformers [1], and a variety of stochastic shortest-path and threshold problems for Markov decision processes [103, 104], all exhibit such hardness.

Petri nets and semi-Markov chains. Termination questions for certain classes of Petri nets are Positivity-hard [6], as is the 'faster than' problem for one-action semi-Markov chains [102].

Probabilistic program equivalence. The universal equivalence of probabilistic programs over finite fields, with applications to differential privacy, has been shown by G. Barthe, C. Jacomme, and S. Kremer [19] to be decidable in important fragments and undecidable in others, with Positivity-style questions arising at the boundary.

Sequential decomposition. In the setting of relational system design, the *sequential decomposition problem* asks whether a task, specified as an input/output relation, can be split into two sequentially composed sub-tasks [55, 56]. For automatic relations the problem is conjectured to be undecidable: already for the equality relation, total decomposability can be viewed as an automata-theoretic problem—given a regular language L, does L contain at most 2^n words of each length n?—which is shown in [56] to be interreducible with the Positivity Problem.

Control theory and matrix semigroups. The Skolem and Positivity Problems appear in a range of decision problems concerning the stability, observability, and controllability of linear control systems; for an entry point to this literature we refer to the survey of V. D. Blondel and J. N. Tsitsiklis [35], and to more recent work on linear time-invariant systems [54] and on matrix semigroups [24].

9.3 Open Problems

We close with a short list of concrete challenges, in the expectation that progress on any one of them would move the whole area forward.

Skolem at order 5. The outstanding configuration at order 5 comprises simple LRS whose dominant roots form two complex-conjugate pairs of common modulus (Section 5.3). Decidability at order 5 holds conditionally on the Exponential Local-Global Principle [76], and for arbitrary simple LRS assuming additionally the weak p-adic Schanuel Conjecture (Theorem 33); a different conditional route, through the strengthened Cramér–Granville Conjecture, is described in Section 6.5. All of these conjectures are expected to hold, but proving any one of them would constitute a major breakthrough in number theory. An unconditional proof of decidability at order 5—or, more ambitiously, for all simple LRS, perhaps beginning with suitably structured sub-families—is in our view the most immediate challenge in the area.

The order-6 sequence of Section 8.3. Establish that the sequence (13) has no natural-number zero, or exhibit one. No modular certificate exists (Proposition 53), no semialgebraic invariant exists, and separation bounds lie beyond current Baker-type technology: settling this single instance therefore seems to require an extension of the existing certificate classes, with potential applicability well beyond the instance itself.

Positivity for simple LRS at order 10. Decide Positivity for simple LRS of order 10—for example, for the concrete sequence (16) of Section 8.4. By Theorem 52, decidability at order 10 would also resolve the Skolem Problem at order 5.

Positivity at order 6. For general LRS, decidability of Positivity at order 6 would, by Theorem 51, yield the computability of Diophantine-approximation constants that have resisted all attempts to date, and is accordingly not expected in the near term. By contrast, no such barrier is known for simple LRS at order 10, which is why the preceding challenge appears more tractable.

Universal Skolem Sets. The known constructions (Sections 6.4 and 6.5) produce Universal Skolem Sets of density one (in the case of the set $\mathcal{S}_1$ of Section 6.4, conditional on the Bateman–Horn Conjecture). Exhibiting substantially richer universal sets—ideally with sparse, structurally simple complements—is wide open; the endpoint of this line of research, showing that $\mathbb{N}$ itself is a Universal Skolem Set, is the assertion that the Skolem Problem is decidable.

Complexity at low orders. The Skolem Problem for LRS of order at most 4 lies in coRP (Theorem 26); whether the randomness can be eliminated is open. (The known NP-hardness applies only to LRS of unbounded order, and so does not stand in the way.)

The contemporary picture that emerges from this survey is one of substantial recent activity—unconditional advances at low orders alongside a wealth of conditional results—against a backdrop of stubborn open questions whose resolution will demand new advances in number theory, new algorithmic ideas in the certificate framework, or both.

Acknowledgements

P. Bacik and J. Worrell are supported by EPSRC grant EP/X033813/1. P. Bacik and J. Ouaknine are supported by ERC grant DynAMiCs (101167561) and DFG grant 389792660 as part of TRR 248. J. Nieuwveld is supported by the Glasstone Benefaction, University of Oxford [Violette and Samuel Glasstone Research Fellowships in Science 2025]. J. Ouaknine is also affiliated with Keble College, Oxford as emmy.network Fellow.

References

[1] R. Aghamov, C. Baier, T. Karimov, J. Nieuwveld, J. Ouaknine, J. Piribauer, and M. Vahanwala. Model checking Markov chains as distribution transformers. In *Principles of Verification (2)*, volume 15261 of *Lecture Notes in Computer Science*, pages 293–313. Springer, 2024.

[2] S. Akshay, T. Antonopoulos, J. Ouaknine, and J. Worrell. Reachability problems for Markov chains. *Information Processing Letters*, 115(2):155–158, 2015.

[3] S. Akshay, N. Balaji, A. Murhekar, R. Varma, and N. Vyas. Near-optimal complexity bounds for fragments of the Skolem problem. In *Proceedings of the International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 154 of *LIPIcs*, pages 37:1–37:18. Schloss Dagstuhl, 2020.

[4] S. Akshay, N. Balaji, and N. Vyas. Complexity of restricted variants of Skolem and related problems. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 83 of *LIPIcs*, pages 78:1–78:14. Schloss Dagstuhl, 2017.

[5] S. Akshay, H. Bazille, B. Genest, and M. Vahanwala. On robustness for the Skolem, positivity and ultimate positivity problems. *Logical Methods in Computer Science*, 20(2), 2024. Article 11.

[6] S. Akshay, S. Chakraborty, A. Das, V. Jagannath, and S. Sandeep. On Petri nets with hierarchical special arcs. In *Proceedings of the International Conference on Concurrency Theory (CONCUR)*, volume 85 of *LIPIcs*, pages 40:1–40:17. Schloss Dagstuhl, 2017.

[7] P. B. Allen. On the multiplicity of linear recurrence sequences. *Journal of Number Theory*, 126:212–216, 2007.

[8] E. Allender, P. Bürgisser, J. Kjeldgaard-Pedersen, and P. B. Miltersen. On the complexity of numerical analysis. *SIAM Journal on Computing*, 38(5):1987–2006, 2009.

[9] S. Almagor, D. Chistikov, J. Ouaknine, and J. Worrell. O-minimal invariants for discrete-time dynamical systems. *ACM Transactions on Computational Logic*, 23(2):7:1–7:20, 2022.

[10] F. Amoroso and E. Viada. On the zeros of linear recurrence sequences. *Acta Arithmetica*, 147:387–396, 2011.

[11] P. Bacik. Completing the picture for the Skolem problem on order-4 linear recurrence sequences. *TheoretiCS*, 4(28), 2025.

[12] P. Bacik, J. Ouaknine, D. Purser, and J. Worrell. On the p-adic Skolem problem. In *Proceedings of the International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 364 of *LIPIcs*, pages 8:1–8:20. Schloss Dagstuhl, 2026. Extended version, correcting a technical lemma, at arXiv:2504.14413.

[13] P. Bacik, J. Ouaknine, and J. Worrell. On the complexity of the Skolem problem at low orders. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 5255–5269. SIAM, 2026.

[14] P. Bacik and A. Varonka. On the subspace orbit problem and the simultaneous Skolem problem. In *Proceedings of the Annual Symposium on Logic in Computer Science (LICS)*, volume 380 of *LIPIcs*, pages 8:1–8:27. Schloss Dagstuhl, 2026.

[15] C. Baier, F. Funke, S. Jantsch, T. Karimov, E. Lefaucheux, F. Luca, J. Ouaknine, D. Purser, M. A. Whiteland, and J. Worrell. The orbit problem for parametric linear dynamical systems. In *Proceedings of the International Conference on Concurrency Theory (CONCUR)*, volume 203 of *LIPIcs*, pages 28:1–28:17. Schloss Dagstuhl, 2021.

[16] C. Baier, F. Funke, S. Jantsch, T. Karimov, E. Lefaucheux, J. Ouaknine, A. Pouly, D. Purser, and M. A. Whiteland. Reachability in dynamical systems with rounding. In *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 182 of *LIPIcs*, pages 36:1–36:17. Schloss Dagstuhl, 2020.

[17] A. Baker. *Transcendental Number Theory*. Cambridge University Press, 1975.

[18] A. Baker and G. Wüstholz. *Logarithmic Forms and Diophantine Geometry*. New Mathematical Monographs. Cambridge University Press, 2007.

[19] G. Barthe, C. Jacomme, and S. Kremer. Universal equivalence and majority of probabilistic programs over finite fields. *ACM Transactions on Computational Logic*, 23(1):5:1–5:42, 2022.

[20] B. Bartolomé, Y. Bilu, and F. Luca. On the exponential local-global principle. *Acta Arithmetica*, 159(2):101–111, 2013.

[21] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in real algebraic geometry*. Springer, 2006.

[22] J. P. Bell and S. Gerhold. On the positivity set of a linear recurrence sequence. *Israel Journal of Mathematics*, 157:333–345, 2007.

[23] P. C. Bell, J.-C. Delvenne, R. M. Jungers, and V. D. Blondel. The continuous Skolem-Pisot problem. *Theoretical Computer Science*, 411(40–42):3625–3634, 2010.

[24] P. C. Bell, I. Potapov, and P. Semukhin. On the mortality problem: From multiplicative matrix equations to linear recurrence sequences and beyond. *Information and Computation*, 281:104736, 2021.

[25] J. Berstel and M. Mignotte. Deux propriétés décidables des suites récurrentes linéaires. *Bulletin de la Société Mathématique de France*, 104:175–184, 1976.

[26] J. Berstel and C. Reutenauer. *Noncommutative Rational Series with Applications*. Cambridge University Press, 2010.

[27] C. Bertók and L. Hajdu. A Hasse-type principle for exponential Diophantine equations and its applications. *Mathematics of Computation*, 85(298):849–860, 2016.

[28] C. Bertók and L. Hajdu. A Hasse-type principle for exponential Diophantine equations over number fields and its applications. *Monatshefte für Mathematik*, 187(3):425–436, 2018.

[29] F. Beukers and R. Tijdeman. On the multiplicities of binary complex recurrences. *Compositio Mathematica*, 51(2):193–213, 1984.

[30] Y. Bilu. Skolem problem for linear recurrence sequences with 4 dominant roots (after Mignotte, Shorey, Tijdeman, Vereshchagin and Bacik), 2025. arXiv:2501.16290.

[31] Y. Bilu, F. Luca, J. Nieuwveld, J. Ouaknine, D. Purser, and J. Worrell. Skolem meets Schanuel. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 241 of *LIPIcs*, pages 62:1–62:15. Schloss Dagstuhl, 2022.

[32] Y. Bilu, F. Luca, J. Nieuwveld, J. Ouaknine, and J. Worrell. On the p-adic zeros of the Tribonacci sequence. *Mathematics of Computation*, 93(347):1333–1353, 2023.

[33] Y. Bilu, F. Luca, J. Nieuwveld, J. Ouaknine, and J. Worrell. Twisted rational zeros of linear recurrence sequences. *Journal of the London Mathematical Society*, 111(3):e70126, 2025. arXiv:2401.06537.

[34] V. D. Blondel and N. Portier. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear Algebra and its Applications*, 351–352:91–98, 2002.

[35] V. D. Blondel and J. N. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.

[36] A. Brumer. On the units of algebraic number fields. *Mathematika*, 14:121–124, 1967.

[37] J. R. Burke and W. A. Webb. Asymptotic behavior of linear recurrences. *Fibonacci Quarterly*, 19(4):318–321, 1981.

[38] J.-Y. Cai, R. J. Lipton, and Y. Zalcstein. The complexity of the ABC problem. *SIAM Journal on Computing*, 29(6):1878–1888, 2000.

[39] F. Calegari and B. Mazur. Nearly ordinary Galois deformations over arbitrary number fields. *Journal of the Institute of Mathematics of Jussieu*, 8(1):99–177, 2009.

[40] J. W. S. Cassels. *Local Fields*, volume 3 of *London Mathematical Society Student Texts*. Cambridge University Press, 1986.

[41] V. Chonev, J. Ouaknine, and J. Worrell. The polyhedron-hitting problem. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 940–956. SIAM, 2015.

[42] V. Chonev, J. Ouaknine, and J. Worrell. On recurrent reachability for continuous linear dynamical systems. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 515–524. ACM, 2016.

[43] V. Chonev, J. Ouaknine, and J. Worrell. On the complexity of the orbit problem. *Journal of the ACM*, 63(3):23:1–23:18, 2016.

[44] V. Chonev, J. Ouaknine, and J. Worrell. On the Skolem problem for continuous linear dynamical systems. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 55 of *LIPIcs*, pages 100:1–100:13. Schloss Dagstuhl, 2016.

[45] V. Chonev, J. Ouaknine, and J. Worrell. On the zeros of exponential polynomials. *Journal of the ACM*, 70(4):26:1–26:26, 2023.

[46] H. Cohen. *A Course in Computational Algebraic Number Theory*, volume 138 of *Graduate Texts in Mathematics*. Springer, 1993.

[47] J. D'Costa, T. Karimov, R. Majumdar, J. Ouaknine, M. Salamati, S. Soudjani, and J. Worrell. The pseudo-Skolem problem is decidable. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 202 of *LIPIcs*, pages 34:1–34:21. Schloss Dagstuhl, 2021.

[48] J. D'Costa, T. Karimov, R. Majumdar, J. Ouaknine, M. Salamati, and J. Worrell. The pseudo-reachability problem for diagonalisable linear dynamical systems. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 241 of *LIPIcs*, pages 40:1–40:13. Schloss Dagstuhl, 2022.

[49] H. Derksen. A Skolem–Mahler–Lech theorem in positive characteristic and finite automata. *Inventiones Mathematicae*, 168(1):175–224, 2007.

[50] H. Derksen and D. Masser. Linear equations over multiplicative groups, recurrences, and mixing II. *Indagationes Mathematicae*, 26(1):113–136, 2015.

[51] R. Dong and O. Shafrir. The Skolem problem in rings of positive characteristic. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*. ACM, 2026. To appear.

[52] G. Everest, A. van der Poorten, I. Shparlinski, and T. Ward. *Recurrence Sequences*, volume 104 of *Mathematical Surveys and Monographs*. American Mathematical Society, 2003.

[53] J.-H. Evertse. On sums of S-units and linear recurrences. *Compositio Mathematica*, 53(2):225–244, 1984.

[54] N. Fijalkow, J. Ouaknine, A. Pouly, J. Sousa Pinto, and J. Worrell. On the decidability of reachability in linear time-invariant systems. In *Proceedings of the ACM International Conference on Hybrid Systems: Computation and Control (HSCC)*, pages 77–86. ACM, 2019.

[55] D. Fried, A. Legay, J. Ouaknine, and M. Y. Vardi. Sequential relational decomposition. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 432–441. ACM, 2018.

[56] D. Fried, A. Legay, J. Ouaknine, and M. Y. Vardi. Sequential relational decomposition. *Logical Methods in Computer Science*, 18(1):37:1–37:29, 2022.

[57] V. Halava, T. Harju, and M. Hirvensalo. Positivity of second order linear recurrent sequences. *Discrete Applied Mathematics*, 154(3):447–451, 2006.

[58] V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem's problem — on the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, 2005.

[59] E. Hrushovski, J. Ouaknine, A. Pouly, and J. Worrell. Polynomial invariants for affine programs. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 530–539. ACM, 2018.

[60] E. Hrushovski, J. Ouaknine, A. Pouly, and J. Worrell. On strongest algebraic program invariants. *Journal of the ACM*, 70(5):29:1–29:22, 2023.

[61] R. Kannan and R. J. Lipton. Polynomial-time algorithm for the orbit problem. *Journal of the ACM*, 33(4):808–821, 1986.

[62] T. Karimov. *Algorithmic Verification of Linear Dynamical Systems*. PhD thesis, Saarland University, 2024.

[63] T. Karimov, E. Kelmendi, J. Nieuwveld, J. Ouaknine, and J. Worrell. The power of positivity. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–11. IEEE, 2023.

[64] T. Karimov, E. Kelmendi, J. Ouaknine, and J. Worrell. What's decidable about discrete linear dynamical systems? In *Principles of Systems Design*, volume 13660 of *Lecture Notes in Computer Science*, pages 21–38. Springer, 2022.

[65] T. Karimov, E. Lefaucheux, J. Ouaknine, D. Purser, A. Varonka, M. A. Whiteland, and J. Worrell. What's decidable about linear loops? *Proceedings of the ACM on Programming Languages*, 6(POPL):1–25, 2022.

[66] M. Kauers and P. Paule. *The Concrete Tetrahedron: Symbolic Sums, Recurrence Equations, Generating Functions, Asymptotic Estimates*. Texts and Monographs in Symbolic Computation. Springer, 2011.

[67] E. Kelmendi. Computing the density of the positivity set for linear recurrence sequences. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 3:1–3:12. ACM, 2022.

[68] G. Kenison, J. Nieuwveld, J. Ouaknine, and J. Worrell. Positivity problems for reversible linear recurrence sequences. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 261 of *LIPIcs*, pages 130:1–130:17. Schloss Dagstuhl, 2023.

[69] N. Koblitz. *p-adic Numbers, p-adic Analysis, and Zeta-Functions*, volume 58 of *Graduate Texts in Mathematics*. Springer, 2nd edition, 1984.

[70] J. C. Lagarias and J. O. Shallit. Linear fractional transformations of continued fractions with bounded partial quotients. *Journal de Théorie des Nombres de Bordeaux*, 9(2):267–279, 1997.

[71] S. Lang. *Algebraic Number Theory*, volume 110 of *Graduate Texts in Mathematics*. Springer, 2nd edition, 1994.

[72] V. Laohakosol and P. Tangsupphathawat. Positivity of third order linear recurrence sequences. *Discrete Applied Mathematics*, 157(15):3239–3248, 2009.

[73] C. Lech. A note on recurring series. *Arkiv för Matematik*, 2:417–421, 1953.

[74] E. Lefaucheux, J. Ouaknine, D. Purser, and M. Sharifi. Model checking linear dynamical systems under floating-point rounding. In *Proceedings of the International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 13993 of *Lecture Notes in Computer Science*, pages 47–65. Springer, 2023.

[75] A. Lindenmayer and G. Rozenberg, editors. *Automata, Languages, Development*. North-Holland, 1976.

[76] R. J. Lipton, F. Luca, J. Nieuwveld, J. Ouaknine, D. Purser, and J. Worrell. On the Skolem problem and the Skolem conjecture. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 5:1–5:9. ACM, 2022.

[77] R. J. Lipton and K. W. Regan. *People, Problems, and Proofs — Essays from Gödel's Lost Letter: 2010*. Springer, 2013.

[78] F. Luca, J. Maynard, A. Noubissie, J. Ouaknine, and J. Worrell. Skolem meets Bateman–Horn. *Computing Research Repository (CoRR)*, abs/2308.01152, 2023.

[79] F. Luca, J. Ouaknine, and J. Worrell. Universal Skolem sets. In *Proceedings of the Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–6. IEEE, 2021.

[80] F. Luca, J. Ouaknine, and J. Worrell. Algebraic model checking for discrete linear dynamical systems. In *Proceedings of the International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 13465 of *Lecture Notes in Computer Science*, pages 3–15. Springer, 2022.

[81] F. Luca, J. Ouaknine, and J. Worrell. A universal Skolem set of positive lower density. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 241 of *LIPIcs*, pages 73:1–73:12. Schloss Dagstuhl, 2022.

[82] F. Luca, J. Ouaknine, and J. Worrell. On large zeros of linear recurrence sequences. In *Proceedings of the International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 345 of *LIPIcs*, pages 71:1–71:11. Schloss Dagstuhl, 2025.

[83] F. Luca, J. Ouaknine, and J. Worrell. Conjectural decidability of the Skolem problem. *Computing Research Repository (CoRR)*, abs/2607.15510, 2026.

[84] K. Mahler. Eine arithmetische Eigenschaft der Taylor-Koeffizienten rationaler Funktionen. *Proceedings of the Akademie van Wetenschappen, Amsterdam*, 38:50–60, 1935.

[85] R. Majumdar, J. Ouaknine, A. Pouly, and J. Worrell. Algebraic invariants for linear hybrid automata. In *Proceedings of the International Conference on Concurrency Theory (CONCUR)*, volume 171 of *LIPIcs*, pages 32:1–32:17. Schloss Dagstuhl, 2020.

[86] N. Mariaule. *p-adic exponential ring, p-adic Schanuel's conjecture and decidability*. PhD thesis, University of Manchester, 2014.

[87] D. Marques and T. Lengyel. The 2-adic order of the Tribonacci numbers and the equation $T_n = m!$. *Journal of Integer Sequences*, 17, 2014. Article 14.10.1.

[88] D. W. Masser. Linear relations on algebraic groups. In A. Baker, editor, *New Advances in Transcendence Theory*, pages 248–262. Cambridge University Press, 1988.

[89] E. M. Matveev. An explicit lower bound for a homogeneous rational linear form in logarithms of algebraic numbers, II. *Izvestiya: Mathematics*, 64(6):1217–1269, 2000.

[90] M. Mignotte. Calculs sur des suites récurrentes linéaires. *Mémoire de la SMF*, 49:50, 1977.

[91] M. Mignotte, T. N. Shorey, and R. Tijdeman. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik*, 349:63–76, 1984.

[92] J. Müllner, M. Moosbrugger, and L. Kovács. Strong invariants are hard: On the hardness of strongest polynomial invariants for (probabilistic) programs. *Proceedings of the ACM on Programming Languages*, 8(POPL):882–910, 2024.

[93] K. Nagasaka and J.-S. Shiue. Asymptotic positiveness of linear recurrence sequences. *Fibonacci Quarterly*, 28(4):340–346, 1990.

[94] J. Neukirch. *Algebraic Number Theory*, volume 322 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1999.

[95] J. Nieuwveld. *Algorithmic Problems for Linear Recurrence Sequences*. PhD thesis, Saarland University, 2025. Work carried out at the Max Planck Institute for Software Systems (MPI-SWS), Germany.

[96] J. Ouaknine. Holonomic techniques, periods, and decision problems (invited talk). In *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 182 of *LIPIcs*, pages 4:1–4:3. Schloss Dagstuhl, 2020.

[97] J. Ouaknine, J. Sousa Pinto, and J. Worrell. On termination of integer linear loops. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 957–969. SIAM, 2015.

[98] J. Ouaknine and J. Worrell. On the positivity problem for simple linear recurrence sequences. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 8573 of *Lecture Notes in Computer Science*, pages 318–329. Springer, 2014.

[99] J. Ouaknine and J. Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 366–379. SIAM, 2014.

[100] J. Ouaknine and J. Worrell. Ultimate positivity is decidable for simple linear recurrence sequences. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 8573 of *Lecture Notes in Computer Science*, pages 330–341. Springer, 2014.

[101] J. Ouaknine and J. Worrell. On linear recurrence sequences and loop termination. *ACM SIGLOG News*, 2(2):4–13, 2015.

[102] M. R. Pedersen, N. Fijalkow, G. Bacci, K. G. Larsen, and R. Mardare. Timed comparisons of semi-Markov processes. In *Proceedings of the International Conference on Language and Automata Theory and Applications (LATA)*, volume 10792 of *Lecture Notes in Computer Science*, pages 271–283. Springer, 2018.

[103] J. Piribauer and C. Baier. On Skolem-hardness and saturation points in Markov decision processes. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 168 of *LIPIcs*, pages 138:1–138:17. Schloss Dagstuhl, 2020.

[104] J. Piribauer and C. Baier. Positivity-hardness results on Markov decision processes. *TheoretiCS*, 3(9), 2024.

[105] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. Part I. *Journal of Symbolic Computation*, 13(3):255–299, 1992.

[106] A. M. Robert. *A Course in p-adic Analysis*, volume 198 of *Graduate Texts in Mathematics*. Springer, 2000.

[107] G. Rozenberg and A. Salomaa. *Cornerstones of Undecidability*. Prentice Hall, 1994.

[108] A. Salomaa and M. Soittola. *Automata-theoretic Aspects of Formal Power Series*. Texts and Monographs in Computer Science. Springer, 1978.

[109] A. Schinzel. On primitive prime factors of $a^n - b^n$. *Proceedings of the Cambridge Philosophical Society*, 58:555–562, 1962.

[110] A. Schinzel. Abelian binomials, power residues and exponential congruences. *Acta Arithmetica*, 32(3):245–274, 1977.

[111] A. Schinzel. On the congruence $u_n \equiv c \pmod{p}$, where u_n is a recurring sequence of the second order. *Acta Academiae Paedagogicae Agriensis, Sectio Mathematicae*, 30:147–165, 2003.

[112] W. M. Schmidt. The zero multiplicity of linear recurrence sequences. *Acta Mathematica*, 182(2):243–282, 1999.

[113] A. Schönhage. On the power of random access machines. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 71 of *Lecture Notes in Computer Science*, pages 520–529. Springer, 1979.

[114] M. P. Schützenberger. On the definition of a family of automata. *Information and Control*, 4(2–3):245–270, 1961.

[115] T. Skolem. Ein Verfahren zur Behandlung gewisser exponentialer Gleichungen und diophantischer Gleichungen. In *Comptes Rendus du Huitième Congrès des Mathématiciens Scandinaves*, pages 163–188, 1934.

[116] T. Skolem. Anwendung exponentieller Kongruenzen zum Beweis der Unlösbarkeit gewisser diophantischer Gleichungen. In *Avhandlinger Norske Videnskaps-Akademi i Oslo I, No. 12*, pages 1–16, 1937.

[117] M. Soittola. On D0L synthesis problem. In A. Lindenmayer and G. Rozenberg, editors, *Automata, Languages, Development*. North-Holland, 1976.

[118] R. P. Stanley. *Enumerative Combinatorics, Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2nd edition, 2012.

[119] L. J. Stockmeyer and A. R. Meyer. Word problems requiring exponential time (preliminary report). In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, pages 1–9. ACM, 1973.

[120] T. Tao. *Structure and Randomness: Pages from Year One of a Mathematical Blog*. American Mathematical Society, 2008.

[121] S. P. Tarasov and M. N. Vyalyi. Orbits of linear maps and regular languages. In *Proceedings of Computer Science — Theory and Applications (CSR)*, volume 6651 of *Lecture Notes in Computer Science*, pages 305–316. Springer, 2011.

[122] M. Vahanwala. Skolem and positivity completeness of ergodic Markov chains. *Information Processing Letters*, 186:106481, 2024.

[123] A. J. van der Poorten and H. P. Schlickewei. Additive relations in fields. *Journal of the Australian Mathematical Society*, 51(1):154–170, 1991.

[124] A. J. van der Poorten and H. P. Schlickewei. Zeros of recurrence sequences. *Bulletin of the Australian Mathematical Society*, 44(2):215–223, 1991.

[125] N. K. Vereshchagin. Occurrence of zero in a linear recursive sequence. *Mathematical Notes of the Academy of Sciences of the USSR*, 38(2):609–615, 1985.

[126] M. Waldschmidt. *Diophantine Approximation on Linear Algebraic Groups: Transcendence Properties of the Exponential Function in Several Variables*, volume 326 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2000.

[127] K. Yu. p-adic logarithmic forms and group varieties II. *Acta Arithmetica*, 89(4):337–378, 1999.